

NASA/CR-97-206258



Semiannual Report

April 1, 1997 through September 30, 1997

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA*

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contracts NAS1-97046 & NAS1-19480

November 1997

Available from the following:

NASA Center for AeroSpace Information (CASI)
800 Elkridge Landing Road
Linthicum Heights, MD 21090-2934
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

CONTENTS

	Page
Introduction	ii
Research in Progress	
Applied and Numerical Mathematics	1
Fluid Mechanics	22
Applied Computer Science	31
Reports and Abstracts	46
ICASE Colloquia	62
ICASE Summer Activities	66
Other Activities	72
ICASE Staff	73

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics;
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period April 1, 1997 through September 30, 1997 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-97046. Financial support was provided by NASA Contract Nos. NAS1-97046, NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

BRIAN ALLAN

Reduced Order Model for Sensor/Actuator Placement Problem

Design of an advanced aircraft incorporating micro sensors and actuators leads to a sensor/actuator placement problem. In order to gain insight into this problem, a model of the Navier-Stokes equations is used in an LQR optimal feedback control setting. The spatial structure of the feedback kernel will identify regions where feedback sensitivity and actuation are large. As a result of the large cost of calculating the optimal feedback kernel, a reduced order model needs to be developed. In this investigation, we will calculate an optimal feedback kernel using a reduced order model of the incompressible Navier-Stokes equations, providing guidance on the placement of sensors and actuators.

The reduced order model is generated by calculating the exponential growth of small disturbances in a steady-state base flow. Krylov methods are used to approximate this exponential operator by finding its largest eigenvalues and eigenmodes. This approximation leads to a reduced order model of the Navier-Stokes equations, linearized about a nominal steady-state flow field. Once a reduced order model has been developed, it will then be incorporated into an LQR optimal feedback control design. To evaluate this methodology, the feedback gains for a two-dimensional shear flow problem are being investigated. A code which solves the unsteady incompressible flow equations, using a vorticity-stream function formulation, has been developed. A reduced order model has also been calculated by estimating the leading eigenmodes of the exponential operator. These eigenmodes are estimated using the ARPACK software package coupled with the flow solver. We are currently evaluating this reduced order model.

Once a reduced order model has been developed, the optimal feedback gains will be calculated providing guidance on the placement of sensors and actuators in the flow field. This investigation will be extended to the flow around an airfoil. This work is done in collaboration with J. Loncaric at ICASE.

EYAL ARIAN

Coupling Theory for Multidisciplinary Analysis and Optimization

The objective of this work is to develop quantitative estimates of how tightly coupled is a given system of coupled PDEs with respect to analysis, sensitivity, and optimization computations. This work is relevant to algorithm development for problems in multidisciplinary analysis and optimization.

The approach is to derive expressions for the coupling based on the partial differential operators involved and then to estimate an upper bound for these expressions. The predictions of the analysis were verified numerically on simplified model problems including a non-linear algebraic system and a system of linear PDEs (static potential flow over a plate).

Future plans are to further investigate the applications of the above methodology to MDO problems.

Optimization Problems for Radome Design

The objective of this work is to determine the optimal thickness and permittivity of a multilayer Radome such that some desired properties are attained. (Radome is the term used for the nose of the aircraft where the antennas are located).

Our approach is to pose these problems as optimization problems with possible inequality constraints on the design variables (thickness and permittivity) and solve them with gradient based algorithms.

We have been able to solve two different problems.

The first problem is to determine the optimal thickness of a symmetric three layer Radome (with thicknesses (d_1, d_2, d_1)), averaging over a specific frequency band (around 89 GHz) and incident angles (around 29°), such that the transmission of radiation through the Radome is maximal.

The second problem is to optimize the thickness d and real and imaginary parts of the permittivity, ϵ_1 and ϵ_2 , of a single layer Radome such that the resulting transmission and reflection of radiation through the Radome match given measurements of that data.

This work was done in collaboration with F.B. Beck, C.R. Cockrell, C.J. Reddy and M.D. Deshpande.

Development of Efficient Algorithms for Large Scale Aerodynamic Optimization

The purpose of this project is to develop algorithms which do not require more than a few full solutions of the flow equations to obtain the optimum.

We have completed the implementation of a new method on optimal shape design of airfoils, governed by the Euler equations, using pointwise discretization of the shape. The new method is based on preconditioning the gradient direction with an approximation of the inverse of the Hessian in the continuum level (Arian and Ta'asan, ICASE Report No. 96-28). We intend to further investigate such methods in order to install a robust and efficient design capability into TLNS3D.

This work is done in collaboration with V. Vatsa of NASA Langley.

H.T. BANKS

Application of Smart Material Technology to Nondestructive Damage Detection in Plate-Like Structure

The focus of this research is to apply smart material technology to the nondestructive detection and characterization of damage in plate-like structures using vibration data from self-sensing, self-actuating structures. Among the different types of smart materials currently studied, structures with bonded piezoelectric ceramic patches are of particular interest. When an electric field is applied, piezoelectric patches induce strains in the materials to which they are bonded and, conversely, they produce a voltage when a deformation occurs in the material. As a consequence, these patches can act both as actuators and sensors, providing the host structure with smart material capabilities.

We have developed a mathematical model for a damaged 2-D plate-like structure with piezoelectric patches bonded onto its surface. Computational methods for the detection of damage and characterization of physical parameters such as stiffness, damping, or mass density are being developed.

Experiments to characterize damaged structures were conducted to verify our computational algorithms with experimental data.

This research is conducted in collaboration with P. R. Emeric (North Carolina State University) and W. P. Winfree (NASA Langley Research Center).

HESTER BIJL

Comparison of Two Numerical Methods for the Computation of Low Mach Number Flow

My topic of research at the Delft University of Technology is the development and implementation of a numerical method which does not suffer from loss of efficiency and accuracy when low Mach number regions are present in compressible flows. A staggered grid pressure correction method which uses scaled variables and a pressure difference unknown is developed for the computation of flows at all speeds. The results obtained for inviscid flows are very good: the convergence rate and accuracy of our method does not deteriorate when low Mach numbers are present. Recently the method is extended to viscous flows, for which similar results were obtained. Also at ICASE work is done into the computation of low Mach number flow. Dr. Turkel, a.o., has worked on a time derivative preconditioning method to solve the efficiency and accuracy problem at low Mach numbers. Time derivative preconditioning methods start from the equations, where we start from the variables, and multiply a term (the time derivative) in the equation by a specially designed matrix. The objective of my project at ICASE was:

- Investigation of the relation of our scaling with the matrix multiplication in preconditioning methods
- Computation of the same test problem with medium and low Mach numbers by both methods and comparison of the results.

The approach was to first do a literature survey and talk to Dr. Turkel to fully understand the preconditioning methods; then to select a test example, preferably one that was already run by Dr. Turkel, and compute this with our method for low Mach number flow. The theoretical comparison is completed, and a test problem, inviscid flow over the NACA0012 airfoil, is selected. The grid of Dr. Turkel is embedded in the code at Delft and the first computations with our low Mach number method using our standard discretisation in general coordinates were carried out. Unfortunately, this discretisation needs smooth grids, and the grid provided was not smooth enough. Therefore, no solution is obtained yet. Computations at zero Mach number with the improved discretisation, developed at Delft and implemented only for incompressible flow, did look good. Now that our methods are compared theoretically I conclude that the scaling in our method can be expressed as a diagonal preconditioning of the entire system, opposed to the preconditioning of the time derivative only in the method of Dr. Turkel. As a result, in time derivative preconditioning the eigenvalues of the system change, while in our case they do not. The aim of scaling the eigenvalues is to obtain a less stiff system. Unfortunately, it destroys time accuracy, so that pseudo time stepping is required for the computation of instationary flows. For low Mach numbers our method reduces to a classical staggered grid pressure correction method. This is an efficient method for time-dependent incompressible flows. To conclude, our methods are very different, only part of both is a -different- matrix multiplication in the Euler equations.

In future the improved discretisation will be implemented for compressible flows. Thereafter the inviscid flow over the NACA0012 airfoil will be computed and efficiency and accuracy of the two methods will be compared for this test problem.

ACHI BRANDT

Advanced Multigrid Solvers in Aerodynamics

The efficiency of current multigrid solvers in aerodynamics lags several orders of magnitude behind the ideal (or “textbook”) efficiency that has been obtained long ago for uniformly elliptic systems. The main

reason for this inefficiency is the inadequate treatment of the various nonelliptic factors of the Euler or Navier-Stokes system of equations, and especially the relatively short distance to which smooth characteristic components are approximated on coarse grids. Other computational obstacles have also failed to be systematically studied and treated. Our objective now has been to first prepare a comprehensive list of all these obstacles, their possible treatments and current states of development, and then to push systematically toward their resolution.

The basic solution approach is to separate out the various factors of the PDE system by the technique of distributive and weighted relaxation, so that each factor can be treated individually, taking precisely into account its ellipticity measures and characteristic directions relative to each grid of the multigrid cycle. Each computational barrier associated with each such factor should then be studied through carefully-chosen model problems. A road map for such developments has been prepared in a form of a table, listing many of the relevant factors and barriers, each with a referenced list of possible solutions and current status. Included in the table are staggered and nonstaggered, conservative and nonconservative discretizations of viscous and inviscid, incompressible and compressible flows at various Mach numbers, with some simple (algebraic) turbulence modelling. The listing of associated computational barriers involves: non alignment of streamlines and sonic characteristics with the grids; recirculating flows; stagnation points; discretization and relaxation on and near shocks and boundaries; far-field artificial boundary conditions; important features (e.g., complete boundaries) which are not visible on the coarse grids; large grid aspect ratios; boundary layer resolution; grid adaption and unstructured grids. A well-organized attack on the multigrid slowness caused by each of these barriers has thus been outlined.

In the future we plan to collaborate with NASA scientists and ICASE visitors on a sequence of model problems, through which the methods to overcome these various barriers can be developed and demonstrated.

The work so far has been in collaboration with Dinar, Diskin, Mavriplis, Melson, Roberts, Ruge, Sidilkover, South, Swanson, Taasan and Thomas.

Predictor-Corrector Technique in CFD Problems

This work was done in conjunction with Boris Diskin, also of the Weizmann Institute of Science, Israel.

The approach we follow to develop a textbook efficient multigrid solver for the Euler system of equations is to use the distributive relaxation technique to separate the treatment of different factors appearing in the principal part of the Euler operator. These factors are the advection operator and the full potential operator. Discretizations of these factors are derived from a discretization of the Euler system. Frequently these discrete factors do not allow solving by a straightforward marching since their discretizations include some forward points and/or the CFL stability condition is broken. We are referring, for example, to a highly accurate discrete advection operator and to the discrete full potential Operator derived from a staggered discretization of the Euler system. Usually in such cases the defect correction technique is used. However, in the case of non-elliptic operators its convergence is too slow. Our objective was to develop efficient predictor-corrector type algorithms for the aforementioned discrete operators and to understand general principles of devising such algorithms for solving non-elliptic operators.

During our one-month visit we focused on understanding inherent disadvantages of the defect correction technique and on correcting them in the framework of predictor-corrector solvers. As a test we have developed an extremely efficient solver for a third order accurate discretization of the 2D advection operator. In our numerical tests the predictor-corrector cycle, requiring just 4 minimal work units, reduces the error by factor 1/7 in the worst case (usually, it is much better). The corresponding FMG algorithm yields a solution

with the discretization accuracy in just 5 minimal work units. We have also experimented with the discrete full potential operator in the deep supersonic regime ($M^2 > 2$, where M is the Mach number). We have developed and tested an efficient predictor-corrector solver for this problem. Numerical experiments showed that this solver, when combined with the transonic full potential operator solver developed earlier gives us a textbook-efficient multigrid solver for the full potential operator at all Mach numbers.

Our main goal remains to solve the Euler system with textbook multigrid efficiency. Therefore our efforts will be mainly directed to translate the existing solvers for the Euler system factors into a distributive relaxation scheme for the entire Euler system discretized in primitive variables. Another objective related to the predictor-corrector approach is the development of a stable efficient solver for unstable discretizations, such as a discrete wave equation violating the CFL stability condition.

DAVID DARMOFAL

Enhancing Robustness and Efficiency of Multigrid Algorithms for the Navier-Stokes Equations

In recent years, CFD has become an integral tool in the design of aerospace vehicles. However, Navier-Stokes codes remain plagued by significant demands on processing times and memory sizes. Thus, a renewed interest has developed in increasing the performance of Navier-Stokes methods, in particular, accelerating the convergence of these algorithms. A method which has the potential to significantly decrease computational time is the combination of multigrid with local preconditioning. Unfortunately, many local preconditioners, especially those designed for good low Mach number performance, are known to have significant robustness problems especially at stagnation points. The objective of this research is to develop improved local preconditioners which not only accelerate convergence but also are robust and reliable.

A major contributor to robustness problems with local preconditioners is that the eigenvectors of the system degenerate as the Mach number approaches zero. This degeneracy permits significant amplification of perturbations in the flow and thus leads to a lack of robustness especially at stagnation points. Thus, the approach taken in this research is to develop techniques to limit this amplification. Specifically, we developed a limiting technique which is based on local flow gradients. This technique guarantees that when local flow perturbations are high, the possible amplification is low, thus bounding the overall size of a perturbation. Similarly, in regions where the flow is smooth, the local limiter is not active and larger amplifications are allowed. Previous limiting strategies utilized reference Mach numbers and required problem dependent tuning. Next, this new, local limiting strategy was tested in a several multigrid codes including a variety of inviscid and viscous simulations using both structured and unstructured grids. Mach number ranges included nearly incompressible flows to transonic flows with shocks. The results obtained thus far have shown the new limiting strategy to be very robust.

Future plans for this work involve completing additional test cases to further validate the soundness of the approach. Finally, a complete report of the approach and results will be written for ICASE and journal publication in the next few months.

Note: During the course of this work at ICASE, the author had many helpful conversations on preconditioning, robustness, and multigrid with several people, especially: Dimitri Mavriplis and Eli Turkel. Both have also helped tremendously in the numerical testing of the limiting approach using their multigrid codes.

GUY GENDRON

Thermoelastic Analysis of Laminated Plates Using a $\{1-2\}$ -Order Theory. Finite Element Implementation

Organic-matrix composites offer superior structural characteristics due to their high stiffness and strength-to-weight ratios. As a result, they are seen as good material candidates for the next generation of supersonic airplanes. However, their low out-of-plane resistance requires that all the components of the stress tensor be taken into account during the design process. This poses a great challenge to classical structural analysis procedures for plate and shell structures. For example, in the case of plate structures, the widely used Reissner/Mindlin kinematic assumption results in constant values of the out-of-plane shear components. To more accurately predict the structural behavior of composite material structures, a new plate theory has been formulated. This so-called $\{1,2\}$ -order plate theory is based on in-plane displacement fields that vary linearly through the thickness and a quadratic expression for the out-of-plane displacement field. The theory has recently been extended to include thermal loading cases similar to those that will be seen during normal flight conditions of supersonic planes. The purpose of the work was to derive the consistent load vector corresponding to the $\{1,2\}$ -order plate theory for the case of thermal loadings.

This work has been performed in the NASA Computational Mechanics Testbed software (COMET). This general-purpose finite element software already implements a triangular element based on the $\{1,2\}$ -order plate theory. The capabilities of the element have been extended so that a through the thickness linearly-varying temperature field can now be applied to the finite element model. When such a loading case is applied, the element calculates the corresponding consistent nodal load vectors. These vectors are then assembled to calculate the displacements that will result from the application of the temperature field.

The future work will include the development of a test suite that will more extensively verify the implementation of the new functionality. This suite will include thermal loading cases for isotropic and laminated plates. It will also include combined loading cases consisting of temperature gradients and pressure loads. Also, the functionality currently only allows for the calculation of the displacement field. It must be extended to include the calculation of the stress field resulting from the application of the thermal load case.

The author wishes to thank Dr. Alexander Tessler and Mrs. Christine A. Lotts for their help in this effort.

OMAR GHATTAS

PDE-Constrained Optimization on Parallel Computers: Algorithms and Applications

I spent two weeks at ICASE as a participant in the 1997 Summer Visitor Program. My primary host was D. Keyes; I was a visitor in the area of Applied and Numerical Mathematics. My specific interest was in parallel solution of large scale optimization problems arising in optimal design and inverse problems. My main objective was to interact, share ideas, and explore opportunities for collaboration with (1) D. Keyes on parallel "optimization-friendly" preconditioners for unstructured mesh problems; (2) the larger set of ICASE researchers interested in optimal design and optimal control, including M. Lewis and E. Arian; and (3) more generally others at ICASE interested in large-scale scientific computation and visualization.

My main activity during my visit was discussions with various researchers. D. Keyes helped me identify and formulate a class of unstructured mesh preconditioners that were memory-efficient, reasonably effective, and appropriate for multiple right hand sides (as required by sensitivity analysis), and that fit into the

parallel framework I was using. These have been implemented into my parallel optimization code and are being currently tested. I discussed common areas of interest with the following ICASE staff and visitors (area in parentheses): M. Lewis (optimization, inverse problems), E. Arian (optimization, coupled multidisciplinary problems), B. Allan and J. Loncaric (optimal control), T. Crockett and K. Ma (visualization, parallel rendering), A. Pothen (sparse matrix methods, graph partitioning), S. Guattery (graph partitioning, preconditioners), and P. Forsyth (PDE solvers). I also interacted with others at NASA Langley with common interests, including K. Anderson and a number of researchers in the MDO Branch (P. Newman, J. Barthelmy, L. Green, N. Alexandrov, G. Hou, and C. Gumbert).

Work will continue on implementing optimization-friendly parallel preconditioners, which will greatly reduce run times and allow solution of more realistic problems. Plans have been made for the CMU Earthquake Modeling Grand Challenge group to use the ICASE parallel unstructured volume renderer when it becomes available for the Cray T3E. This will allow us to do on-line rendering, thereby significantly increasing the resolution of our earthquake wave propagation visualizations, which have been limited by the sequential rendering tools we've used.

DAVID GOTTLIEB

Computational Methods in Electromagnetics

The objective effort involves development of multi-domain high-order schemes for electromagnetics.

A general multi-domain spectral code for the computation of scattering from two-dimensional objects of arbitrary cross-section such as cylinders has been completed and tested. The developed software is found to be significantly faster than FD-TD (Finite Difference Time Domain) schemes while yielding superior accuracy as measured in the bistatic radar-cross-section. The code is able to handle broadband excitation with the same ease as it handles one frequency!

Completed and tested an implementation of body-of-revolution scattering code. Special attention has been given to the treatment of coordinate singularities and edge boundary conditions. Through the multi-domain approach, completely arbitrary axis-symmetric objects can be handled and the code has been successfully tested for objects such as spheres, cone-spheres and finite cylinders, yielding accurate results, as compared with e.g. MoM (Methods-of-Moments) solutions, at a significantly reduced computational effort. As a more general test, we have computed scattering from a missile to find acceptable agreement with MoM solutions. This code is particularly well suited for high-frequency and transient problems and can handle incident waves at arbitrary incidence. However, the new code handles with the same ease the case of many frequencies.

We have shown that the PML split-form of Maxwell's equations is only *weakly well-posed* and therefore its solution could diverge under small perturbations. We also proposed a different set of equations that seem to offer some advantages. The properties of its solutions for the same geometries and boundary condition were also investigated.

Work has been done in conjunction with S. Abarbanel (Tel-Aviv University, Israel), J. Hesthaven (Brown University), and B. Yang.

JAN HESTHAVEN

The Analysis and Construction of Wellposed and Stable PML Method for the Equations of Acoustics

One of the most important, and as yet generally unsolved, problems associated with the solution of wave-dominated problems is the often encountered need for introducing an artificial exterior boundary at which boundary conditions needs to be specified. A few years back it was suggested that non-physical sponge layers, now known as PML-layers, be used and such schemes were developed for electromagnetics and recently also for the equations of acoustics. While such layers were shown to be perfectly absorbing irrespective on the angle of incidence and frequency of the incoming waves, severe stability problems were encountered in the resulting numerical schemes. It has been the purpose of the present study to arrive at an understanding of the source of the observed instability and attempt a reformulation of the PML scheme that avoids such problems.

A complete analysis of the original PML scheme for the equations of acoustic has been completed, providing the explanation of the numerical instability in terms of loss of strong wellposedness and illposedness under low order perturbations. A new strongly wellposed PML scheme for non-convecting free stream acoustic problem has been derived and analyzed and an alternative sponge layer method, exhibiting PML-like behavior, is proposed. The behavior and stability of the scheme, as expected from the analysis, has been confirmed through experiments. This analysis and subsequent developments of wellposed and stable PML schemes provides a mathematically sound alternative to the previously developed PML schemes for the equations of acoustics.

We shall continued the development of wellposed and stable PML methods for the general case of convecting free-stream acoustics.

LELAND JAMESON

Adaptive Multidomain Methods using Wavelets

This project was intended to combine many ideas to arrive at a numerical method which adapts not only the density of grid points but also the order of the method using wavelets. Furthermore, the method is load balanced for efficient parallel implementation. The objective is a practical and efficient load balanced wavelet-based adaptive numerical method.

This project required only a few months to complete and to date has shown great promise on solving hyperbolic problems efficiently in the areas of computational aeroacoustics and computational electrodynamics.

In the future we plan to adapt the method to nontrivial geometries and to apply it to real-world problems in computational electrodynamics.

This work was done in collaboration with Jan Hesthaven (Brown University).

R. MICHAEL LEWIS

Managing Approximations and Surrogates in Optimization

The enormous computational cost of complex scientific and engineering simulations makes it impractical to rely exclusively on high-fidelity simulations for the purpose of design optimization. Instead, we will need to make as much use as possible of models of lower physical fidelity but lower computational cost, with only

occasional recourse to expensive, high-fidelity simulations. Such use of approximation models is in keeping with engineering practice, where models of lower fidelity are widely used in preliminary design to explore the design space. We are investigating the systematic and algorithmic use of such approximations and surrogates in the context of optimization.

We are pursuing an approach based on the trust region idea from nonlinear programming that provides a systematic response to both poor and incorrect prediction on the part of the approximation, while not being so conservative as to retard progress when the approximations do a good job of predicting the behavior of the true system. By monitoring how well the approximations are predicting the behavior of the system, the trust region mechanism also suggests guidelines for changing the approximation based on its predictive abilities. We have determined the minimal theoretical requirements that approximations must satisfy for use in optimization algorithms, and have identified large classes of approximations satisfying these conditions.

As part of this work, R. M. Lewis organized an ICASE workshop on surrogates this summer, which brought together a group of roughly a dozen researchers from academia, industry, and NASA to discuss experience with using approximations and surrogates and to identify significant research areas in the field. A report of the week's discussion is in preparation.

We are currently focussing on the development of practical algorithms. We intend to investigate the use of first-order kriging methods as general approximations, together with the use of local/global searches in optimization that attempt to take longer steps by using non-local models. We have also begun an investigation of the use of approximations based on mesh refinement and a posteriori error estimates. We will also examine the use of problem-specific approximations, and the incorporation of information from models of many different levels of physical fidelity. We have identified a number of potential test cases, including structural optimization of a wing-box, helicopter rotorblade design, and the design of the Aerospike rocket engine nozzle for the Reusable Launch Vehicle.

This work was done in collaboration with Natalia Alexandrov (NASA Langley Research Center, Multidisciplinary Design Optimizations Branch), and Virginia Torczon (the College of William and Mary and ICASE).

Pattern Search Methods for Optimization

Pattern search methods remain popular with users of optimization because of their simplicity and surprising effectiveness in the unconstrained case. Our recent work in this area has focused on extending pattern search to linearly constrained problems as the part of a general program of treating general nonlinear programs.

We have developed a family of provably convergent pattern search algorithms for linearly constrained minimization. The key is to insure that the patterns have a sufficient density of directions relative to the polar of the cone determined by the normals of the almost binding constraints. This allows us to avoid the problem of premature convergence. This is the first extension of pattern search algorithms to linearly constrained problems.

We have already begun numerical implementation of pattern search algorithms for bound and linearly constrained minimization. With this software we will be able to investigate a number of practical questions concerning the efficacy of this class of algorithms. We expect to transfer this work to V. J. Torczon's collaborators at Boeing.

We also intend to continue our work on generalizing pattern search methods to general nonlinearly constrained optimization. We have developed a precise understanding of the types of geometrical and

analytical problems that we must treat if we are to develop provably robust algorithms. In particular, if some sensitivity information about the constraints is available we should be able to use an exact non-smooth penalty approach.

This work was done in collaboration with Virginia Torczon (the College of William and Mary and ICASE).

Qualitative Nature of Hessians for Problems Associated with PDE

In optimization of systems governed by partial differential equations, the nature of the governing PDE affects the nature of the Hessian of the objective and constraints. At the same time, accurate approximation of second derivatives greatly improves the efficacy of nonlinear programming algorithms. This research focuses on examining the relationship between the nature of the governing PDE and the nature of the Hessian.

Our approach is based on the connection between first and second derivatives in such problems to the reduced gradient and reduced Hessian. In particular, this relation reveals how the solution operator for the PDE enters into the structure of the Hessian operator.

For hyperbolic problems, the Hessian can exhibit hyperbolic behavior insofar as the second-order effects of a perturbation in the optimization variables can propagate significantly to other parts of the region of spatial definition (a la wave motion or propagation of singularities). We have established this phenomenon in the case of travel-time tomography and one-way wave propagation.

The structure of the Hessian is such that if the governing PDE is elliptic or parabolic, the associated Hessian will be something like a pseudodifferential operator, while if the governing PDE is hyperbolic, the Hessian may be either something like a Fourier integral operator or a pseudodifferential operator.

We intend to pursue the analytical aspects of the relationship between Hessians and PDE, and to explore the practical, numerical consequences of this relationship.

JOSIP LONCARIC

The Pseudo-Inverse of the Derivative Operator

Spectral factorization of the Laplacian in cylindrical geometry leads to operators of the form $D \pm |k|I$, where D is the derivative operator with respect to $\log(r)$ and k is the circumferential wavenumber. In polynomial spectral methods, these operators are similar to large Jordan blocks which are invertible for $k \neq 0$. The condition number of the similarity transform is $2^{N-1} N!$ where N is the order of the highest retained Chebyshev polynomial. For large N , obtaining the exact inverse is not computationally useful because the similarity transform is effectively singular in machine precision. The commonly used Lanczos tau method resolves the difficulty at the cost of adding a full row to the matrix, even though this inversion problem can be written in terms of a tridiagonal matrix. Our objective was to retain accuracy while keeping the computational cost low.

Integration preconditioners for differential operators in spectral methods are tridiagonal for arbitrary classical orthogonal polynomial families. This work proposed and investigated the use of the *optimal* integration preconditioner: the pseudo-inverse Q of the derivative operator D , defined in terms of orthogonality of polynomials. The operator Q , which behaves similarly to the Fourier integration operator, is a simple tridiagonal matrix whose eigenvalues lie on the imaginary axis. This was proven analytically for a particularly simple family of basis polynomials and conjectured on the basis of numerical and analytic evidence for

Chebyshev or Legendre polynomials. Analytic expressions for eigenvectors as functions of eigenvalues were obtained, and action of Q fully described including the Jordan block associated with the zero eigenvalue. The spectral norm of the preconditioned inverse operator $(I \pm |k|Q)^{-1}Q$ was shown to be small for any k . In Chebyshev representation, we developed a pair of simple resolution criteria which show that N should grow proportionately to $\sqrt{|k|+1}$ in order to maintain the residue smaller than 2.22×10^{-16} . The method is exact (apart from the truncation error) when $k = 0$ or when the right hand side is orthogonal to the constant basis polynomial. Nearly machine accuracy of the proposed method was numerically verified against the exact analytic solutions. The method is also $O(N)$ times faster than the Lanczos tau method.

This work will be used in developing a fast spectral method for solving a control problem for exterior flows. This method is not limited to operators of the form $D \pm |k|I$. One can take advantage of this fast and accurate method by factorization or repeated preconditioning of more general differential operators.

Spatial Structure of an Optimal Exterior Flow Control

Designing distributed control systems begins with the sensor/actuator placement problem. While in some situations discrete search of combinatorial complexity seems unavoidable, continuum problems suggest solving a related question. *If one could sense everything and actuate everywhere, what should one do?* The answer to this question has polynomial complexity (of order N^3 where N is the number of state variables) and can serve as the initial effectiveness filter capable of rejecting a large portion of the design search space. This favorable situation can have several causes. In our investigation we focus on the effect of no-slip boundary conditions on an optimal flow control of the unsteady Stokes flow around a cylinder. This test problem aims at the development of numerical methods capable of solving the problem of stabilizing the desired flow around wings at low to moderate Reynolds numbers.

Our approach begins by defining a *pseudodifferential representation* ξ of the flow. This step, similar to the Wiener-Hopf factorization of the Laplacian, leads to an explicit diagonalization of the system dynamics $\xi_t = \mathcal{A}\xi$ by means of the Fourier and Weber transforms. We then pose and solve an optimal distributed LQR problem with gain ϵ . The optimal feedback kernel is shown to be well approximated by the resolvent $\epsilon^2(2\mathcal{A} - \epsilon I)^{-1}$ which we compute analytically. For the region where $\epsilon R \ll 1$, we derive the asymptotic form of this nearly optimal feedback in terms of vorticity and demonstrate that boundary actuation and distributed sensing (or estimation) dominate at low circumferential wavenumbers, while at high wavenumbers and far from the boundary colocated sensing/actuation is dominant. Furthermore, we test a very accurate spectral numerical scheme and demonstrate that while some modes are modelled correctly to machine accuracy, nonphysical behavior is introduced as soon as the discretization fails to satisfy the Nyquist criterion for the relevant eigenfunction.

This analytic solution demonstrates the importance of careful spatial discretization such that *all* eigenfunctions of interest are simultaneously resolved throughout the domain. In order to determine viscosity imposed resolution requirements, we intend to extend this approach to Stokes flows around the NACA 0015 wing. Once the viscosity imposed discretization requirements are clear, we intend to investigate control of low and moderate Reynolds number flows.

LI-SHI LUO

Method of Lattice Boltzmann Equation for Complex Fluids

Hydrodynamics of complex systems such multi-phase or multi-component flow are interesting, yet these systems are difficult to simulate by using conventional numerical techniques such as finite difference, etc.

The method of lattice Boltzmann equation (LBE) is one of promising technique for complex fluids. The LBE method has a number of advantages over conventional methods: (1) parallelism of LBE method; (2) simplicity of programming; and (3) ease of incorporating model interactions. Our objective is to establish a consistent theory of thermodynamics of LBE method for complex fluids.

We start with a system of continuum Boltzmann equation, with all thermodynamic quantities well defined. We follow the discretization procedure proposed by He and Luo to derive a LBE for complex fluids. By using this approach, we have gained insights on the discrepancy between the continuous Boltzmann equation and its corresponding lattice Boltzmann equation, and on systematic improvements on the LBE algorithm.

We have worked out the thermodynamics of multi-phase systems (in collaboration with Dr. Hudong Chen of EXA Corporation). We have completed direct numerical simulations of cavity flow by LBE methods. We have also found out the connection between LBE method and beam scheme (in collaboration with Prof. Kun Xu of HKUST).

We will continue the work on complex fluids, in both theoretical study and direct numerical simulation.

DIMITRI MAVRIPLIS

Unstructured Multigrid Convergence Acceleration for Highly Stretched Meshes

For high-Reynolds number viscous flow simulations, efficient resolution of the thin boundary layer and wake regions requires mesh spacings several orders of magnitude smaller in the normal direction than in the streamwise direction. This extreme grid stretching results in poor multigrid convergence rates, usually one to two orders of magnitude slower than those observed for equivalent inviscid flow problems without grid stretching. The purpose of this work is to devise improved multigrid techniques for acceleration convergence on anisotropic grids of this type.

The approach taken consists of using implicit line solvers in the direction normal to the grid stretching combined with semi-coarsening or directional coarsening multigrid to alleviate the stiffness due to grid anisotropy. A graph algorithm has been implemented which combines edges of the original mesh into lines which follow the direction of maximum coupling in the unstructured mesh. A similar algorithm is used to selectively coarsen the original fine mesh by removing points along the directions of strong coupling, thus recursively generating a hierarchical set of coarse meshes for the multigrid algorithm. These coarse levels are constructed using the method of agglomeration. While the original isotropic agglomeration method was based on an unweighted-graph algorithm, the new directional agglomeration procedure relies on a weighted-graph algorithm. Preconditioning techniques to alleviate low-Mach number induced stiffness in the equations have also been implemented. Finally, a Newton GMRES technique has been implemented which uses the above described multigrid and preconditioning techniques as a preconditioner themselves. Using this approach, convergence rates for viscous flows which are independent of the degree of mesh stretching can be obtained. In two-dimensions, this results in an order of magnitude increase in efficiency in the asymptotic convergence range, and about a factor of four increase for engineering calculations.

The multigrid strategy described above is currently being implemented in three dimensions where the resulting code is to be applied to three-dimensional high-lift computations. For large cases, this will require a distributed memory parallel implementation which will involve research into partitioning unstructured grids taking into consideration the directional features of the solver.

Recently, the scientific community has expressed a renewed interest in the analysis of short-pulse radiation and transient scattering phenomena. The characterization of transient wave phenomena is of paramount importance in disciplines ranging from electromagnetics to acoustics, elastodynamics, and geophysics. Efficient computational analysis of these phenomena hinges upon the availability of fast time domain algorithms. Today, all prevailing time domain techniques for analyzing electromagnetic wave scattering phenomena are differential-equation based. Examples include the finite difference time domain technique and the time domain finite element method. Unfortunately, the computational complexity of these solvers scales unfavorably as the electrical dimensions of the object become large. For free-space surface radiation and scattering problems, Time Domain Integral-Equation (TDIE) based methods appear to be preferable to differential-equation based techniques because (i) they only require a discretization of the scatterer surface, (ii) they implicitly impose the radiation condition, and (iii) they are devoid of grid dispersion errors. Unfortunately, TDIE techniques have long been conceived as intrinsically unstable and computationally expensive when compared to their differential-equation counterparts. In this work, we attempt to develop TDIE solvers which rely on novel multilevel paradigms with reduced computational complexity when compared to classical TDIE schemes.

We have developed novel Plane Wave Time Domain (PWTD) algorithms which accelerate the computational analysis of transient surface scattering phenomena using TDIE techniques. The PWTD algorithms developed thus far permit the fast evaluation of transient fields satisfying the wave equation. The cost associated with the computation of fields at N_s observers produced by a surface bound source density represented in terms of N_s spatial samples for N_t time steps scales as $O(N_t N_s^2)$ if classical time domain integral-equation based methods are used. We have shown that this cost can be reduced to $O(N_t N_s^{4/3})$ and $O(N_t N_s \log(N_s))$ using two-level and multilevel PWTD schemes, respectively. These algorithms are the time domain counterparts of frequency domain fast multipole methods and make feasible the practical broadband analysis of scattering from large and complex bodies. It has been shown that the error in the observer fields depends solely on the approximations introduced by spherical and temporal interpolation functions, whose construction parameters can be chosen in accordance with any desired error criterion.

In the future, we will apply two-level and multilevel PWTD algorithms to the analysis of large-scale acoustic and electromagnetic scattering problems. We will also attempt to extend the PWTD algorithms to multilayered media and to hybridize the PWTD algorithm with shooting and bouncing ray methods. If successful, we anticipate that our algorithms will provide the engineering community with efficient tools for analyzing the interaction of transient acoustic and electromagnetic waves with structures of unprecedented electrical dimensions and geometrical complexity.

Collaborators: A. Arif Ergin, Balasubramaniam Shanker, Center for Computational Electromagnetics, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801.

FREDERIC NATAF

A Krylov-Schwarz method with Ventcell Interface Conditions and Nonconforming Grids for the Helmholtz Equation

My proposal consisted in working on extending the Optimized Order 2 method (OO2) to non matching grids using PETSc. We (Francois Rogier, Francis Nier, Caroline Japhet, Philippe Chevalier and Frederic Nataf) have developed new interface conditions involving second-order derivatives in the tangential direction, whose coefficients are optimized with respect to the speed of convergence of Krylov methods. More precisely, compared to the classical Krylov-Schwarz method, we use these boundary conditions in place of the Dirichlet boundary conditions. The improvement is very significant. This work was done for conforming grids on the interface.

The success of the mortar method (Bernardi, Maday and Patera) has created interest of being able to consider nonconforming (non matching) grids on the interface. It is therefore tempting to use the OO2 method in conjunction with the mortar method. This is difficult on the face of it, since the mortar method makes use of Dirichlet/Neumann boundary conditions. Hence, a new approach to nonconforming grids had to be developed. My proposal was to work on this for the Helmholtz equation. I considered the Helmholtz equation since it is seemingly the scalar equation for which the problem is the most difficult. This equation is not dissipative so that numerical errors generated on the nonconforming interface don't vanish away inside the subdomains. On the other hand, a very simple finite volume discretization was chosen such as not to make things too intricate. Discretization on the interface was very carefully written. I wrote a finite volume code. In its present form, it is a C++ code for only two nonoverlapping subdomains in 2D and Robin boundary conditions which are simpler than the Ventcell boundary conditions. From preliminary tests, the quality of the solution does not seem to be affected by the nonconformity of the grids, though I use more general boundary conditions than Dirichlet/Neumann boundary conditions.

This first step is very important in that it shows the possibility to use general interface conditions on nonconforming grids. I intend to finish this work by adding support for the Ventcell conditions actually used in the OO2 method. I'll also add the possibility to consider general nonoverlapping domain decompositions. The interface will include corners which demand a specific treatment. I intend to publish my first report on this work as an ICASE report.

Though this was not a joint work with an ICASE member, I had discussions with David Keyes which were very fruitful. It gave us the following idea which is related to the OO2 method (but not to the nonconformity of the grid). The PETSc code developed at Argonne has been widely used in CFD. David Keyes played a role in promoting its use in CFD through the Newton-Krylov-Schwarz method (NKS). The present form of the Krylov-Schwarz implementation in PETSc is based on Dirichlet/Neumann boundary conditions. The idea would be to add the possibility to use Optimized Order 2 boundary conditions in place of Dirichlet boundary conditions in a "black box" form for the end-user. Linear solvers would be faster. L. MacInnes (a co-developer of PETSc) seems interested in this project. I hope that through ICASE (and Keyes), a long term collaboration can be established in order to achieve the integration of the OO2 method in PETSc.

Otherwise, I had interesting discussions with D. Gottlieb (Brown University), E. Turkel (Tel-Aviv University), and A. Baggag (University of Minnesota) on domain decomposition methods and with Josip Loncaric on the Dirichlet to Neumann operator.

ROY A. NICOLAIDES

Fast Multiplication Techniques for Computational Electromagnetics Investigator

Discretization of integral equation formulations of Maxwell's equations results in dense complex systems of linear equations of very large order. In many cases, the matrices are of so large that they cannot be stored in memory or be solved by elimination in any reasonable time frame. A number of algorithms, all involving additional approximations, have been proposed to mitigate these scale effects. The best known approach is the Fast Multipole Method. One aspect of the FMM is that it prescribes how the integral equations are discretized. As a result, it is not very practical to use it in conjunction with existing standard electromagnetics codes.

In this project, a new technique is being developed which can be incorporated into existing electromagnetic codes. It is based on recursive processing of the blocks of the coefficient matrix into a form where fast multiplications become possible. Several approaches to this are feasible depending on the smoothness of the kernel of the integral equations. In every case, once the appropriate decompositions are performed the idea is to use (possibly preconditioned) iterative solvers such as GMRES.

Development of these techniques is in a preliminary stage and there is considerable work remaining to be done. In particular, both theory and implementations will require additional development. However, enough information is presently available for basic implementations to be performed and those are currently underway. Following that, implementations for more complex scattering problems are planned.

RONALD NOWACZYK

Engineering Design Team Process

The engineering design process can be enhanced by a better understanding of the team dynamics occurring among engineers and scientists working on engineering design teams. This research is directed toward identifying team principles and processes that are associated with successful engineering design teams. A primary objective is to develop a metric for assessing engineering design team performance.

The research is based on psychological and sociological theories of team behavior. Most of the principles have been based on studies of management teams. While, engineering design teams share a number of features with management teams, their focus on specific products that can be evaluated objectively from a number of different disciplines and perspectives make them different from many management teams. Two empirical studies have been conducted. The first study focused on identifying the factors crucial to team success. The results from surveying approximately 50 engineers led to the development of a team performance scale. That scale was analyzed in the second study. NASA engineers and scientists provided data on team performance from over 40 engineering design teams. The data revealed that the scale was highly reliable and was a valid discriminator between successful and unsuccessful engineering design teams.

The team performance scale is currently being validated against the ratings of team sponsors and team leaders in an effort to generalize the validity of the scale beyond team member perceptions. A web-based version of the team performance scale is under development. The intent is to have a web-based scale that can be used not only for evaluation but also as a diagnostic tool. Profiles of team performance will be compared with those of successful and unsuccessful teams. In those cases, where problems are noted, interventions to improve team performance will be suggested.

CORD ROSSOW

Investigations of a Least Squares Approximation for the Solution of the Navier-Stokes Equations

During the last years, some effort was spent on using Least Squares approximations for the discretization of the Euler- and Navier-Stokes equations. These methods deviate from the classical Finite Volume approach in such that no metric quantities like cell areas or cell volumes are used. Instead, a Least Squares fit is performed over a number of points in order to approximate the derivatives. In the literature available so far, this approach was termed "gridless", indicating that potentially no computational grid as in classical discretization methods might be necessary. This perspective makes the method very attractive, since it is becoming more and more clear that the main obstacle for CFD methods to be routinely employed is the generation of computational meshes for complex, industrially relevant configurations. In the case that just arbitrarily orientated clouds of points could be used, the process of computing the flow around complex configurations would be substantially alleviated. Up to now, the Least Squares method has only been used for the solution of the Euler equations and to solve for viscous laminar flows. No computations of turbulent flows at high Reynolds numbers have been reported so far. Furthermore, no detailed comparison with classical Finite Volume methods has been made to check the consistency and the conservation properties of the method. Therefore, the objectives of the performed research were to explore the potential of the Least Squares approximation for CFD-applications. Especially, the research was focussed on the computation of high Reynolds number turbulent flows and on the comparison with a classical Finite Volume discretization.

The approach taken to achieve the objectives stated above was the following: The basis code used in this research is a two-dimensional cell centered Finite Volume code for the solution of the Euler- and Navier-Stokes-equations. The code uses structured meshes and employs central differencing with artificial dissipation for the discretization of the governing equations. All components involving the Finite Volume formulation (derivatives, time step, source terms, scaling of dissipation etc.) were successively replaced by the Least Squares approximations. In order to obtain the LS-discretization stencil, the same neighbouring cells as in the FV-scheme were used. For the points to enter the LS-approximation, three possibilities were investigated: First, the cell centers themselves; second, the arithmetic average of the cell center of the cell for which the discretization is performed and the corresponding neighbouring cell centers; and third, the midpoints of the common face of neighbouring cells and discretization cell. For the last two possibilities, an arithmetic average of flow quantities was performed, and those averages with the corresponding locations entered the approximation. Solving the Euler-equations, for transonic airfoil flows the differences between these three approaches were quite small, either with regard to shock location or with regard to total pressure losses. Also the comparison to the FV-scheme showed no significant differences. Especially when using the third approach in the LS-discretization, the results between this LS-discretization and the FV-discretization were indistinguishable. A mesh refinement study showed an accuracy of second order for the LS-discretization, accordingly to the FV-scheme. For the solution of the turbulent Navier-Stokes equations the picture changed: With possibilities one and two for the selection of points entering the LS-discretization, the algorithm diverged immediately. Using the third possibility, i.e. the midpoints of corresponding cell faces, the algorithm worked in the same way as the FV-code, with respect to solution quality as well as with respect to convergence properties. The reason for this behaviour was found to be the occurrence of high aspect ratio cells needed to resolve the thin boundary layer. Summarizing, the work performed so far indicates that the points to be selected to enter the Least Squares discretization may not be chosen arbitrarily. The results obtained showed that the method converges towards the correct weak solution, despite the fact that conservation could not

strictly be proven. However, a comparison to tensor notation showed a significant correspondence to the formulation of derivatives and metric quantities in generalized coordinates. Thus, it was concluded that the method may be viewed as a generalized Finite Difference method. The research showed that the Least Squares approach is able to deliver the same accuracy of results as obtained with a Finite Volume approach.

In future work on this topic, it seems necessary to further investigate the correspondence to tensor notation in general coordinates. This may lead to a more rigorous statement concerning the conservation properties. The performed research showed evidence of this, but formal proof is needed. Furthermore, the potential of the method needs to be explored. Here it was shown that results of an FV-discretization can be produced, but the question remains whether anything may be gained beyond. Therefore, the method should be employed on meshes with arbitrarily shaped cells, i.e. not only with quadrilaterals or triangles. Since no metrics are involved but only points, the method may offer a larger flexibility for such situations. The last issue relevant to be investigated is the implementation of upwind schemes. In the work done so far only central differencing was used. To exploit the benefits of high accuracy schemes, upwind discretizations have to be incorporated, and since in the LS-method no fluxes normal to cell faces are computed, other strategies for implementation need to be sought.

DAVID SIDILKOVER

A Simple Essentially Optimally Multigrid Algorithm for the Inviscid Flow Equations

This is an on-going project. The main objective of this work is to develop a simple essentially optimal multigrid solver for the steady Euler equations (both incompressible and compressible), i.e. a solver whose efficiency is similar to one of a solver for Poisson or Full-Potential equations.

A simple discretization of the Euler equations (both compressible and incompressible) that facilitates such a solver was constructed recently. The elliptic factor separates in the form of Poisson (incompressible case) or Full-Potential (compressible case) operators acting on the pressure. It was demonstrated that the optimal multigrid efficiency can be obtained for some model problems using this approach. The algorithm was implemented in the framework of engineering codes both in the context of unstructured triangular and structured Cartesian grids. The essentially optimal multigrid efficiency was already demonstrated for incompressible Euler equations in some complex geometry cases including the flow past a cylinder and Von Karman - Trefftz airfoil.

The future plans include addressing of some even more realistic testcases and to generalizing of the implementation for Navier-Stokes equations.

Multidimensional Upwinding and Textbook Multigrid Efficiency

This is an on-going project. A genuinely multidimensional high-resolution discretization for the steady compressible Euler equations was constructed recently. One of the fundamental advantage of this approach is that the genuinely multidimensional high-resolution mechanism (unlike the standard one) does not damage the stability properties of the scheme. The next main objective is to construct an optimally efficient multigrid solver that relies on the genuinely multidimensional upwind discretizations. The main advantage of this approach is its generality – it is expected to be extendible to the cases/regimes other than steady inviscid subsonic flow.

It is well known that for the advection dominated problems the coarse grid provides only a fraction of needed correction for certain error components. A known cure for this problem is to separate the treatment

of the advection part from the rest of the system. A simple way to achieve this is through the separate discretization of different factors. A more general approach is to discretize the equations in some primitive form but to design a special relaxation (of the distributive Gauss-Seidel type) that distinguishes between the different co-factors of the system. A necessary condition for this, however, is that the discretization must be *factorizable*, i.e. the determinant of the corresponding matrix of discrete operators should be a product of approximations to different co-factors of the system. A recent study has revealed that some genuinely multidimensional schemes are *factorizable*. A particular *factorizable* scheme for the Euler equations has been constructed such that the corresponding discretization of the Full-Potential factor in the case of grid-aligned flow is identical to the Murmann & Cole scheme.

The current work concentrates on constructing an optimally efficient multigrid solver for the subsonic regime. Also possibilities of extending the genuinely multidimensional approach to unsteady flow problems are under investigation.

Towards Genuinely Multidimensional Schemes/Solvers for Unsteady Problems

The previously developed genuinely multidimensional approach towards discretizing the flow equations addressed the steady case: the resulting schemes are second order accurate at the steady-state only. Extending this approach to unsteady problems is expected to provide for the design of very efficient implicit solvers.

It was necessary to revisit the one-dimensional case first and to design a new one-dimensional high-resolution scheme that is compatible with the genuinely multidimensional approach. The resulting scheme becomes implicit if the Courant number is larger than one and is quasi-explicit (i.e. can be dealt with as explicit) otherwise. As the first stage a quasi-explicit case was studied in details. It has been proven that the scheme satisfies the TVD property. An extensive set of numerical experiments was conducted to verify the accuracy and good discontinuity capturing capabilities of the discretization.

The next step will be to consider the implicit case and to construct a fast solver. The latter should rely on what has been done previously for the steady case. This study is expected to result in very efficient solvers for unsteady problems. In particular this approach should provide some very interesting possibilities for the low speed flow.

BAMBANG I. SOEMARWOTO

The Variational Method for Aerodynamic Optimization using the Navier-Stokes Equations

The variational method is an efficient method for computing the design sensitivities in aerodynamic optimization. This method requires the solution of a set of adjoint equations for the computation of the design sensitivities of the aerodynamic functionals with respect to the shape. The Navier-Stokes equations represent a complete modelling of the compressible viscous flow. A correct formulation of the adjoint equations and the design sensitivities for the Navier-Stokes equations will enable industrial applications of the aerodynamic optimization methodology.

This research aims at obtaining the correct formulation of the adjoint equations and the expression for the design sensitivities (the gradient). These are implemented together with an existing two-dimensional flow solver (HI-TASK developed by NLR) and optimization routine (FSQP developed by the University of Maryland) to form an aerodynamic optimization code.

At the present time, the aerodynamic optimization code has shown to be able to solve a reconstruction-type inverse problem of an airfoil in a transonic flow condition. Starting from the NACA 0012 airfoil, the

target pressure distribution of the RAE 2822 airfoil can be closely recovered within 17 flow analyses, where each flow analysis corresponds to a distinct solution of the Navier-Stokes equations. The computational results strongly indicate that the adjoint equations and the gradient have been formulated correctly.

The methodology will be extended for dealing with three-dimensional flow and for solving practical problems such as drag-reduction subject to aerodynamic and geometric constraints.

VIRGINIA TORCZON

Optimization for Engineering Design

This work focuses on engineering design optimization problems for which the objective is assumed either to be nondifferentiable or for which the available derivatives are unreliable. Such problems occur frequently in practice for problems defined by computer simulations rather than by analytic functions. Similar difficulties also occur when surrogate functions are used in lieu of expensive computer simulations to ameliorate the cost of the optimization process. Our objective is to design, analyze, and implement practical robust nonlinear programming techniques for such instances.

Currently, we have been working to extend the analysis that has been done for pattern search methods for unconstrained problems to handle problems with constraints. This involves a careful and systematic study of what constitutes a sufficient set of search directions for pattern search methods. We had already discovered the minimal requirements for the unconstrained case and a sufficient set for problems with bound constraints; both results appear as ICASE technical reports. The work during this reporting period has determined the requirements for the minimal set for solving problems with general linear constraints; we are in the process of writing up this result for release as an ICASE technical report. This latest result considerably broadens the class of engineering problems to which these techniques may be applied and points the way toward handling problems with general nonlinear constraints. In addition, we have designed, implemented, and tested a simple pattern search for use with surrogate functions. This approach directly inherits the analysis that has been done for pattern search methods while taking into consideration the complicated question of trying to solve a problem by using less-accurate surrogate information. This preliminary study has recently been released as an ICASE technical report.

We plan to continue our investigation of what it takes to ensure a sufficient set of search directions so that we can further extend our results. In particular, we plan to look at problems for which the objective function is not differentiable. Results in this case will allow us to consider problems with general nonlinear constraints, which can be handled by using simple, nondifferentiable merit functions that penalize infeasibility. Once these results are obtained, we will have a suite of robust pattern search methods for handling all common (continuous) nonlinear programming problems. We also expect to continue to work toward integrating these results with approaches designed to make use of surrogate functions.

This has been joint work with Robert Michael Lewis, ICASE, NASA Langley Research Center and Michael W. Trosset, University of Arizona.

ELI TURKEL

Preconditioning

One of the major difficulties with solving steady state viscous problems is the large amounts of computer time required to achieve a steady state solution. This is especially problematical for low Mach number flows

due to the stiffness of the Euler equations at low Mach numbers. In addition numerous researchers have found that the accuracy of the numerical approximation deteriorates as the Mach number decreases to zero.

We introduce a preconditioning matrix in order to equalize the eigenvalues of the modified system. It can also be shown that this also improves the accuracy of the numerical solution as the Mach number goes to zero. This has been implemented in the TLNS3D code. This was then used to solve low Mach number flows over an ONERA wing and also a two dimensional 3 element airfoil. The improvement from the preconditioning both in terms of convergence rate and accuracy was appreciable.

Future plans include making the code more robust and extend the convergence acceleration to transonic flows

This work is done together with Veer Vatsa of NASA Langley

High Order Methods

Some difference for a first order hyperbolic system seem to be first order accurate while the same scheme for the equivalent second order equation is higher order. We seek an explanation of this anomaly.

It was found that some schemes for the first order system seem to be equivalent to schemes for the second order wave equation. Nevertheless, the scheme for the first order system is high order only for the phase velocity but not the amplitude. It is possible to improve the situation by pre- and post- processing the solution.

This work was done together with David Gottlieb and Saul Abarbanel

High Order Methods for the Maxwell Equations

The numerical solution of the time dependent Maxwell equations requires enormous computer resources since one is interested in solving problems that have very high frequencies. Due to the large bandwidth desired many mesh points are needed to achieve reasonable accuracy. In addition many time steps are needed to calculate the solution of long periods of time.

We use a compact implicit method based on a staggered grid similar to that used by Yee. This enables a code based on the Yee algorithm to be simply converted to fourth order accuracy. It also allows a PML absorbing layer to be easily used. Some model problems demonstrate the increased efficiency that can be gained by such an approach.

Future plans include extending this work to discontinuous coefficients and to non-Cartesian bodies.

This work was done together with Amir Yefet a doctoral student at Tel Aviv University.

BRAM VAN LEER

Toward Robust and Efficient Stagnation Preconditioning

In spite of recent concentrated efforts to come up with a local Euler preconditioner that remains robust in the presence of stagnating flow, yet does not lose its preconditioning (convergence accelerating) power, the Advanced Subsonic Technology community is still waiting for a satisfactory result in this area. Our previous attempt[†] led to a preconditioning matrix that has optimal eigenvalues and also preserves orthogonality among the eigenvectors of waves moving parallel to the flow. This, however, appears to be insufficient to preserve stability near a stagnation point. The goal of the present research is to find an algebraic criterion for developing preconditioners that trade optimality of eigenvalues for non-degeneracy of eigenvectors.

[†]Collaboration with Dohyung Lee, University of Michigan

There are several different ways of proceeding in this research, which we will explore all. Currently I am assuming a certain form of the preconditioner, which yields non-optimal eigenvalues but is still clearly related to preconditioners developed before, such as Van Leer-Lee-Roe; it has two free parameters. Another approach would be to apply an energy norm to the preconditioned equations, and search for preconditioners that do not lead to unbound growth for $M \rightarrow 0$. A third approach is to interpret the preconditioning as a way to decouple the system of equations while rescaling it, and to come up in particular with a preconditioned set acoustic equations that remain valid in the limit of $M \rightarrow 0$. In all cases the algebra is very complex, and progress is slow, even with the use of Mathematica or Maple. I am still exploring the first research line; currently the other two look more promising.

Once a preconditioner is found that maintains the lowest possible condition number as $M \rightarrow 0$, while limiting transient growth, it must be tested in computational practice, especially in airfoil calculations, where multiple stagnation points may occur. In order to establish its benefit in accelerating convergence, the preconditioner must be combined with single-grid, standard multi-grid and semi-coarsened multi-grid relaxation.

This work was done in conjunction with David Darmofal of Texas A & M University.

FLUID MECHANICS

ALVIN BAYLISS

Jet Acoustics/Structure/Flow Field Interaction for Non-Axisymmetric Jet Disturbances

We numerically simulate the interaction between sound from a jet and an array of curved, flexible aircraft-type panels. The vibration of the panels in response to jet sound excitation results in structural fatigue and increased interior noise levels. In previous work we have considered panel response for a Cartesian jet in 2 dimensions. In our present work we have extended the previous research to account for cylindrical jets and both axisymmetric and nonaxisymmetric jet disturbances. The objectives of the present research are to compute the jet flow field, near field and far field, in particular the panel loading, and panel response and radiation for cylindrical jets subject to both axisymmetric and nonaxisymmetric excitation.

We consider a fully three dimensional, nonlinear model in cylindrical coordinates for the Euler equations for the jet. The panel loading will be computed from the Euler equations and used as input to Nastran to compute the panel response. A linearized Euler code will then be used to compute the panel radiation. The primary obstacle of this research is the extremely large programs and lengthy computation times required for the nonaxisymmetric computations. We have completed coding the 3 dimensional Euler equation program and have tested it for a variety of nonaxisymmetric jet excitations. At the present time we have computed switching modes in the jet, where the cylindrical angle of maximum acoustic radiation from the jet switches in time.

In the future we plan to simulate helical as well as switching modes of jet disturbances and to determine panel response under loading by radiation from such modes. This work is being done jointly with Lucio Maestrello of NASA Langley Research Center.

GREGORY A. BLAISDELL

Compressible turbulence simulation and modeling

The objective of this work is to validate the pseudo-sound compressible turbulence theory of J. R. Ristorcelli using direct numerical simulation (DNS) of compressible homogeneous turbulence and, more generally, to help increase our understanding of compressible turbulence.

Several DNS of decaying isotropic turbulence and homogeneous shear flow were done at various turbulent Mach numbers and Reynolds numbers. The DNS results match the scaling of the pseudo-sound theory at early times but then show a discrepancy. A higher-order perturbation expansion of the pseudo-sound theory was carried out and seems to indicate that the differences between the theory and the DNS results are due to acoustic waves produced from the initial conditions as opposed to acoustic waves generated by the turbulence. In addition, perturbation methods were applied to the problem of rapid distortion theory for compressible homogeneous shear flow. Some initial results compare well with the numerical results of Cambon and coworkers; higher order corrections have proven difficult.

Future work will involve a more complete perturbation analysis for low Mach number turbulence and for rapidly distorted flows.

Elliptic streamline flow

The elliptic streamline flow is a homogeneous turbulent flow that combines the effects of mean flow rotation and strain. Previous comparisons between direct numerical simulation (DNS) results and standard Reynolds stress turbulence models have shown that the models do not correctly predict whether the turbulent kinetic energy will grow or decay. For strong rotation rates the models predict that the turbulence decays while DNS and linear stability analysis shows that the turbulence grows. The objective of the work for this summer is to see whether more advanced turbulence models are capable of capturing the correct behavior in this flow.

An analysis of linear pressure strain models by S. Girimaji indicates that modifying the model coefficients has an effect on whether the models predict growth or decay for the elliptic streamline flow. It may be possible to modify the current class of turbulent models so that they will capture the correct behavior in the elliptic streamline flow without losing accuracy in other flows.

Data from the DNS were given to S. Girimaji and future work will include a comparison with model predictions. Also, further simulations of the elliptic streamline flow and related homogeneous turbulent flows are planned to provide additional data for model comparisons.

The elliptic streamline flow is related to on-going work involving DNS of a turbulent vortex, which is important to the wake hazard problem for commercial aircraft. While at ICASE a collaboration with F. Proctor of NASA LaRC was established which has helped to explain the effects of atmospheric wind shear on aircraft wake vortices.

PAO-LIU CHOW

Vibrational Control of a Nonlinear Elastic Panel

The problem under investigation is concerned with an active contrail of nonlinear panel oscillation induced by the wall pressure fluctuation. The control objective is to stabilize the panel vibration by a vibrational actuator. This work was motivated by recent experimental studies by L. Maestrello at NASA LaRC. These experimental results showed clearly that a vibrational control can be an effective tool in stabilizing a wall-pressure induced panel vibration as well as a boundary-layer flow. The objective of this research is to develop an analytical method for analyzing the vibrational panel control problem, and to provide a theoretical explanation for the experimentally observed results.

Based on the stability theory for periodic motions, the principle of vibrational control is introduced for dynamical systems with infinite-many degrees of freedom. In particular, we are mainly concerned with the vibrational control of an elastic panel vibration modeled by a nonlinear beam equation forced by the wall-pressure fluctuation. First a general method of vibrational control for some nonlinear evaluation equations is developed. This method is then applied to the panel structure problem. The control used consists of a high-frequency parametric vibration and the forcing frequency modulation. It is shown that, for the unstable panel oscillation near a buckled state, such a combined vibrational control can result in a stable periodic motion. For small amplitude vibration, a perturbation method is used to obtain explicit results for vibrational control. We have successfully developed a theory of vibrational control which, when applied to the nonlinear panel model, gives a qualitative explanation for some experimental data. The analytical approach developed in this research can be used to identify the important parameters which may lead to more effective vibrational control in practice.

The method of vibrational control described above can also be applied to other problems. The experimental evidence, as mentioned previously, showed that the boundary-layer flow could be stabilized by an active control: periodic heating and cooling of the wall. But this has not yet been demonstrated analytically. We plan to continue our research in this area by extending the method of vibrational control to cover the case of active boundary control. This generalized approach will then be used to analyze the flow stabilization problem of interest in search for theoretical explanation for experimental observation.

This research was done in collaboration with Lucio Maestrello at NASA Langley Research Center.

WILLIAM CRIMINALE

Disturbances in Boundary Layers, Early Period Dynamics, Receptivity, and Bypass Transition

The laminar boundary layer on a flat plate continues to be one of the most essential prototypical shear flows in fluid mechanics. This flow is important not only because it is a concept that has vast utility in many diverse physical situations, but because it is a flow that can (and has been) probed extensively by means of theory, experiment, and numerical simulation. Still, after nearly a full century that has been spent in such effort, a complete and thorough understanding of the behavior of the flow in the presence of perturbations has not been obtained. Many important questions remain to be answered before progress towards such goals as positive control of the flow or a means of retarding the transition from laminar to a turbulent state can be achieved.

A combination of analytical and numerical techniques have been used to explore completely arbitrary initial-value problems for the boundary layer on a flat plate. First, linearized governing equations are considered. Then, direct numerical simulation is taken using the input data from the linear results. The full transient period is captured from both a temporal and a spatial basis. The first rational demonstration of boundary layer receptivity is made. And, lastly, bypass transition is understood.

The accomplishments of this work provide a means whereby the question of flow control can now be investigated. Moreover, it provides a firm basis structure for direct numerical simulators.

Collaborators: T.L. Jackson, ICASE, D.G. Lasseigne (Old Dominion University), and R.D. Joslin, NASA Langley are the other members involved in this effort.

Other: Two major research papers have been or will be submitted for publication (All four researchers). Cambridge University Press has requested that a monograph be written for publication and this is in preparation (TLJ and WOC).

SHARATH GIRIMAJI

Non-Equilibrium Complex-Deformation Turbulence: Pressure-Strain Modeling

Many turbulent flows of engineering interest are subjected to complex mean deformation and are not close to the ideal 'equilibrium' state. The current pressure-strain correlation models are calibrated in simple laboratory flows and are generally inadequate for complex flows. A new model suitable for complex flows, as well as the classical flows, is developed.

The model development consists of three distinct steps:

- Equilibrium pressure-strain correlation model suitable for flows subject to complex mean deformation is developed using representation theory and a recently derived theorem for elliptic flows subject to large mean flow vorticity.

- Pressure strain correlation model that is suitable for the rapid distortion limit.
- The complex-deformation non-equilibrium pressure strain correlation model is then derived by interpolating between the models of the two extreme limits determined above.

We have been able to demonstrate the validity and the potential of the new modeling procedure. The next step, is to employ this procedure to further improve the model performance in a wider range of turbulent flows.

Non-Equilibrium Turbulence: Algebraic Stress Modeling

Engineering turbulent flow calculations employ, at best, two-equation turbulence models. These simple-to-use models are typically derived from over-simplified physics and, hence, lack the physical accuracy often required. The objective of this work is to incorporate important complex physics into simple-to-use models.

The algebraic Reynolds stress model that is developed in this work is from the slow manifold solution of the differential Reynolds stress equations. The slow manifold solution is close to the exact differential solution after the effects of the initial conditions have faded. This model development is made possible by a newly postulated algebraic approximation of the differential derivative of the Reynolds stress. After the algebraic approximation is made, the pressure strain correlation model developed in the previous work is used in conjunction with representation theory to derive the final model form.

Extensive validation and testing of the algebraic model in engineering flows is the next step.

Analysis and Modeling of Buoyancy Generated Turbulence Using DNS Data

Turbulent temperature flux plays a crucial role in heated/cooled boundary layers and other important applications. Our objective is to use direct numerical simulation (DNS) data of Rayleigh-Bernard convection to develop and validate turbulent temperature flux models.

In this work, we use DNS data to accomplish four important objectives:

1. Evaluate pressure-strain and pressure-temperature gradient correlation models.
2. To study turbulent transport of Reynolds stress and thermal flux.
3. Examine various modeling assumptions made to derive algebraic Reynolds stress models.
4. To develop and verify fully explicit algebraic models.

Develop and validate similar models for mixed and forced convection.

CHESTER GROSCH

Simulation of a Hot Supersonic Jet in a Lobe Mixer

Two NASA programs, High Speed Research (HSR) and Advanced Subsonics Technology (AST) have identified engine noise reduction as an enabling technology. In the HSR program, jet noise is the principle contributor. In the AST program, jet noise is the principle contributor for aircraft in the current fleet, where a goal of 3dB reduction is established for aircraft engines with bypass ratios up to 5. Methods used to reduce jet noise in both programs often utilize concepts that enhance mixing between high and low speed streams. Of these, the most popular methods utilize concepts that introduce streamwise vorticity. One method of generating streamwise vorticity uses tabs on the edge of the jet. Another method is to have the jet exit thru a lobe mixer. The objective of this study is to assess the mixing effectiveness of a lobe mixer and of tabs and to elucidate the physical mechanism responsible for the increased mixing.

A set of numerical calculations are carried out using the compressible time dependent three dimensional Navier-Stokes equations in two different geometries; a jet with a circular cross section and a lobe mixer. The tabs modeled by pairs of counter rotating vortices. Both “necklace” and “trailing” type vortices are simulated by changing the sense of rotation of the model vortices. The results of the calculations show that both the tabs and a lobe mixer can greatly increase the mixing of the jet and the coflow. The basic physical mechanism of the tab induced mixing is that the streamwise vortices transport the hot, higher momentum fluid from the central region of the jet to the colder, lower momentum region of the coflow and vice versa. The initial configuration of “necklace” vortices is generally unstable while that of “trailing” vortices is stable. In the case of the lobe mixer, the mechanism is an instability of the jet (in the flapping mode) generating streamwise vorticity near the walls followed by an instability of the vortices.

The effect of lobe geometry is being studied. Tabs will be added to the edges of the lobes. Longer range, the effect of “soft” walls will be studied.

EHTESHAM HAYDER

Comparisons of Programming Paradigms

Our goal is to examine issues and efficiencies of parallel libraries and parallel languages for scientific computations on a distributed machine.

In this study we considered a model problem representative of low-order discretizations of nonlinear elliptic PDEs and implemented the same algorithm using both paradigms, i.e, parallel library (PETSc) and parallel language (HPF). Our test case was a steady state solid fuel ignition problem, known as the Bratu problem. In this nonlinear problem, heat generation by combustion process is balanced by heat transfer due to conduction.

A summary of our study will be presented in the 4th National Symposium on Large-Scale Analysis and Design on High Performance Computers and Workstations, Williamsburg, VA. This work was done in collaboration with David Keyes (ODU/ICASE) and Piyush Mehrotra (ICASE).

Non-reflecting Boundary Conditions Our objective is to examine the perfectly Matched layer boundary condition to minimize numerical reflections in aeroacoustic computations.

We examined the effectiveness of the PML condition as non-reflecting computational boundaries for the linearized and nonlinear Euler equations. The PML equations were constructed by the operator-splitting of the governing equations and introducing absorption coefficients in each split equation.

A review of our study was presented in the IUTAM Symposium on Computational Methods for Unbounded Domains, Boulder, CO in the past summer. This investigation was being done in collaboration with Fang Hu (ODU/ICASE), M. Y. Hussaini (FSU) and Harold Atkins (NASA Langley).

GEOFFREY LILLEY

Airframe Noise: Background and Supporting Studies

Predictive procedures for airframe noise are being developed. Studies are being conducted in order to identify and understand the most important features of the airframe noise problem.

- Estimates of dominant noise sources using the theories of Crighton, Ffowcs-Williams and Howe have been made.

- Derivation of velocity laws for undercarriage, flap-edge, and slat noise have been obtained.
- Modified Lighthill's acoustic analogy so that source terms are described in quadratic quantities only.
- Noise suppression in owls has been studied as a model for noise suppression in aerodynamic configurations.

Studies as relevant to airframe noise and in support of the Airframe Noise group's project area at Langley are being conducted.

JAMES MARTIN

Vortex Modeling of Airframe Noise Sources

Sound generated at the side edge of an airfoil's flaps is a very important, in some cases the most intense, source of airframe noise. A dominant feature in the side edge flow field is the trailing vortices which form off the edge of a flap. The interaction of these vortices with the neighboring flap surface plays a major role in the observed noise production. As a detailed computation of the fully three-dimensional flow field in the vicinity of an airfoil's flaps is impossible, it is desirable to investigate simplified models of this type flow and these models' acoustic emission. As a first step in this study, we considered a two-dimensional flow model with the trailing vortices represented as point vortices (Hardin and Martin, 1997). We have now begun work on producing and studying an analogous three-dimensional flow model with the side edge vorticity represented through the use of three-dimensional vortex filaments.

Our initial approach has been to model the flap surface as a rectangular slab of infinite extent in the chordwise direction. We again utilize the same background potential flow as in our earlier two-dimensional investigation. The motion of a vortex filament in the presence of the rectangular slab is computed using a combined vortex method and panel method approach. Quadrilateral vortex panels are used to represent the flap surface with their strength chosen at each time step to satisfy the no penetration condition. Nodes comprising the filament centerline are transported according to the influence of the background potential flow, the flap surface panels, and the Biot-Savart law. To make comparisons with our two-dimensional results, a preliminary code has been completed which utilizes periodicity conditions in the chordwise direction. A comparison of the vortex filament trajectories with the trajectories obtained in our two-dimensional study shows good agreement giving us confidence to extend this study to more realistic flap geometries.

As a next step in this investigation, the periodicity conditions implemented in the chordwise direction will be removed and spatially developing trailing vortices will be considered. This research has been done in collaboration with J.C. Hardin (NASA LaRC).

J.R. RISTORCELLI

Compressible Turbulence: Dilatational Covariances

Models for the effects of the compressibility of turbulence are of crucial importance to the estimation of the effects of compressibility as well as the prediction of turbulent compressible flows.

A systematic perturbation procedure has been developed and used to develop models for the accumulated effects of compressibility on the moment evolution equations of compressible turbulence consistent with the present levels of turbulence models in use ($k-\epsilon$). Models for the pressure dilatation ($\langle pd \rangle$) and the dilatational dissipation (ϵ_c) have been derived.

- The scalings predicted by the theory are consistent with DNS (Blaisdell) and EDQNM (Bertoglio).

- The $\langle pd \rangle$ model has been tested against new homogeneous DNS data. Scalings predicted by the theory are consistent with DNS of isotropic and homogeneous shear turbulence.
- The theory has shown that the effects of the pressure-dilatation are **nominal** expect for non equilibrium flows.
- The theory has shown that the dilatation dissipation are, for most aerodynamic flows, negligible.

A theory for the effects of compressibility has been constructed and validated. The major effects of compressibility are not due to the pressure-dilatation or the dilatation dissipation. The major effects of compressibility appear to be due to the pressure-strain correlation. A model for the pressure strain is possible; an extensive amount mathematical manipulations is necessary.

Compressible Turbulence: Initial Conditions

In work done in collaboration with Gregory Blaisdell, it has been shown that inconsistent initial conditions create spurious wave fields leading to compressible data bases not relevant to the practical engineering problem. Such data bases have been used to produce and calibrate turbulence models. Such turbulence models are as arbitrary as the arbitrary initial conditions used to start DNS. Consistent initial conditions for the DNS of compressible turbulence have been derived. The initial condition on the compressible velocity field are related to time and spatial derivatives of the incompressible pressure.

Further explorations of the existence of unexplained oscillations in the homogeneous shear are being conducted. A two timing procedure has indicated that the ostensible spurious oscillations are related to secular terms. New program areas in aeroacoustics have taken precedent.

Rotating Turbulence: The Reynolds Stresses

Prediction of vortical aerodynamic flows have become important for matters of safety and sound and throughput at airports in work done in collaboration with Sharath Girimaji. A new fundamental result for rapidly rotating flows has been found. It has the same relevance and import to flows with rapid mean flow rotation as the Taylor-Proudman theorem has to flows with rapid frame rotation. The new theorem is an asymptotic result in the spirit of the Taylor-Proudman.

The present class of single-point turbulence closures cannot compute an “elliptical” flow. Building the asymptotic results of the new theorem into the standard class of turbulence models now allows prediction of the so-called elliptic flows. The application of these new results to strongly vortical flows is the subject of current research.

Compressible Corrections to Lighthill's Acoustic Analogy

Our objective is to devise models to predict the sound power radiated from turbulent flows using two-equation and higher turbulence models. An understanding of the compressibility of the acoustic source field for high speed flows is necessary.

A singular perturbation methodology in conjunction with a two-timing approach has been used to close the source terms in Lighthill's acoustic analogy for low turbulent Mach number flows: Lighthill's acoustic analogy has been extended to include the compressibility of the fluctuating motions in the source flow.

This allows assessment of the “incompressible” truncation used in applications of Lighthill's analogy.

It also allows use of Lighthill analogy in situations where compressibility, *modulo* thermoacoustic effects, is important. We intend to:

- 1) Extend the compressible form of Lighthill's acoustic analogy to flows with large mean Mach numbers.
- 2) Use the methods of statistical fluid mechanics to develop a predictive scheme for acoustic radiation using as inputs turbulence quantities calculated in single-point turbulence closures.

ROBERT RUBINSTEIN

Time Correlations and the Frequency Spectrum of Sound Radiated by Turbulent Flows

In work done in collaboration with Ye Zhou of ICASE and IBM, the nature of sound radiation by turbulent flow depends on properties of the turbulent time correlations. Even if the energy contained in motions of any given spatial scale is known, say from the Kolmogorov theory, these motions will radiate sound at a frequency proportional to their inverse correlation time; this dependence determines the frequency distribution of acoustic energy. The present work applies theories of turbulent time correlations to compute frequency spectra of sound radiated by isotropic turbulence and by shear turbulence.

We have applied the *sweeping hypothesis* to calculate the total acoustic power radiated by isotropic turbulence. The power spectrum is shown to scale as $\phi^{-4/3}$ for large frequencies and as ϕ^4 for small frequencies. Some numerical simulation data consistent with these scalings are described by Prof. G. Lilley. We have also evaluate the frequency spectrum of the acoustic radiation from weakly sheared turbulence. Previous studies of "shear noise" show that shear causes a contribution to acoustic radiation which depends on second, instead of fourth-order correlations and approximate the second order correlation by its form in isotropic turbulence. The present analysis attempts instead to evaluate the effects of shear on time correlations and to extract the spectrum of acoustic radiation from the shear-dependent correlation function. A complete theory of noise radiation by shear turbulence must include both effects.

We plan to apply the methodology developed here to jet noise and other aeroacoustics problems.

Development of Models for the Destruction Terms in Dissipation Rate Equations

In another collaborative effort with Ye Zhou of ICASE and IBM, we find that dissipation rate transport equation has been widely used in engineering applications. Basic difficulties arise when there are coupled fluctuating fields as in buoyant turbulence, or imposed time scales as in rotating turbulence. In these cases, even the form of the appropriate transport equation is doubtful. Applying the closure theories, we have developed models for the destruction terms in the dissipation rate equations for incompressible turbulence, buoyant turbulence, rotating incompressible turbulence, and rotating buoyant turbulence. The dependence of the coefficient of the destruction term on external agencies emerges naturally.

The present analysis is based instead on Leslie's suggestion that the dissipation rate transport equation be derived from Kraichnan's integrals for inertial range transport power. The dissipation rate equation for rotating turbulence is derived by substituting rotation-dependent field descriptors in Kraichnan's integrals. To illustrate the general method, the phenomenological theory of rotating turbulence of Zhou is applied. Straightforward application the same theory to buoyant turbulence leads to a preliminary dissipation rate equation for rotating buoyant turbulence. Our development provides a general framework for developing dissipation rate equation for turbulence subject to external agencies.

We plan to apply these to engineering applications; and develop model for other flow fields where external agencies play important role.

LU TING

Turbulent Boundary Layer, Acoustic and Structural Interaction

The interaction of panel oscillation with the boundary layer and incident acoustic waves (simulating jet noise) is essential for the prediction and control of panel fatigue and the reduction of cabin noise. We seek experimental and numerical simulations of the interaction problem and formulate simple theoretical model equations for special aspects of the problem.

We have studied the implications of experimental data obtained by Dr. Lucio Maestrello of NASA La RC on the interaction of turbulent boundary layer and panel oscillations with or without an incident pure tone sound. It was found that the peak levels of the power spectra of panel responses and transmitted waves increase by an order of 20dB in the presence of an incident pure-tone sound. When the level of the pure-tone sound exceeds a threshold value, the characteristics of the panel response change from periodic to chaotic. The panel responses at two measuring points at a distance greater than the boundary layer thickness (and the correlation length) have little correlation. Thus we have the evidence of spatiotemporal chaos, a new phenomenon that has attracted attention in the last ten years.

For numerical simulations of the interaction problem, we examined a recent numerical program developed by Dr. A. Frendi for the interaction of flexible panel and turbulent boundary layer in a supersonic stream. We found that the nonlinear panel equation missed a coupling term with the boundary layer wall shear. The missing term is being added to the program and the parameter characterizing its importance identified.

We intend to formulate simple model equations for the simulation of structural-acoustic interaction problem. We shall present a new derivation of our first and higher order effective on-surface conditions for nonlinear panel oscillations accounting for the acoustic loading and damping effects without the solution of the acoustic field. The new derivation shall then enable us to derive the effective conditions on a convex elastic surface interacting with an acoustic field.

We shall add the missing term in the numerical simulations of nonlinear panel oscillations coupled with turbulent boundary layer and acoustic excitation and show the importance of the missing term, the wall shear stress.

This research has been conducted in close collaboration with Dr. Lucio Maestrello of NASA Langley RC. To develop the numerical program for the simulation of the boundary layer/acoustic/structure interaction, active collaboration of Dr. A. Frendi of Analytical Services and Materials, Inc., Hampton, VA is needed.

APPLIED COMPUTER SCIENCE

ABDELKADER BAGGAG

Parallelization of an Unstructured Discontinuous Galerkin Acoustics Code

Unstructured grids are the choice for many geometrically complex problems in external aerodynamics and acoustics, especially in three dimensions and especially for problems requiring adaptive refinement, where it is difficult to automate the grid generation of any alternatives except for Cartesian grids with irregular boundary cells. In recent years, the discontinuous Galerkin approach has been pioneered on unstructured grids by H. L. Atkins of NASA, for its high-order compact accuracy, and applied explicitly to time-dependent problems acoustic scattering. Atkins' code is sequential, object-oriented, and cleanly modular. Given the important role planned for it, and the intense resolution requirements of wave propagation problems, parallelization is essential.

In order to supply the neighboring entity relationships (e.g., determining the cells on either side of a face, or the vertices at either end of an edge) required in Galerkin discretizations in the distributed memory environment, we employed the Parallel Mesh Environment (PME) of C. Ozturan, developed during 1995-96 and refined during the summer of 1997 at ICASE. The primary loops in Atkins' code are over edges, which fall into two C++ classes, sequentially: interior edges and boundary edges. For the parallel case, the class of interface edges was added using less than 1000 lines of new code to portably parallelize the discontinuous Galerkin code, in conjunction with the PME library. The resulting code has been validated for correctness on the SP2 and on workstation networks; and preliminary performance evaluation is encouraging. This is attributable to the high work-per-point ratio in high-order codes, and the infrequency of global synchronization points in explicit methods.

A careful performance evaluation will be undertaken, including modeling and experimentation. A report will be written codifying the program transformations required and how application programmers can target parallelism in their writing of new unstructured codes. Finally, extensions to higher fidelity physical models, and to variable order discontinuous Galerkin techniques will be pursued.

This work was done in collaboration with D. Keyes (ODU and ICASE), H.L. Atkins (NASA LaRC), and C. Ozturan (Bogazici University, Turkey).

SHAHID BOKHARI

Multithreaded Computation of Unstructured Problems

Computations that use unstructured meshes are notoriously difficult to parallelize. Considerable effort has to be expended on partitioning the mesh over processors and writing suitable message passing code to obtain an efficient implementation. The forthcoming generation of multithreaded shared-memory architectures appears to be an excellent candidate for the solution of such unstructured problems. These architectures eliminate the need for partitioning and thus promise good performance with minimal recoding.

This research aims to implement the EUL3D code developed by Dimitri Mavriplis onto the Tera MTA (MultiThreaded Architecture). We wish to evaluate the performance of the Tera compiler in generating efficient code for this problem and measure the speedups obtained. We also plan to investigate the use of certain hardware features unique to the Tera, such as Full/Empty bits, which promise to eliminate the need for expensive barrier synchronization.

At the present time, the Tera's delivery to the San Diego Supercomputing Center has been delayed. We have experimented with the HP/Convex SPP-2000, which is a 256 processor shared memory machine being installed at Caltech, and obtained useful speedups for small numbers of processors with no recoding. The code has been rewritten using the thread library available on the SPP-2000, however the speedups obtained in this case are unremarkable.

We plan to move the code to the Tera as soon as it becomes available. The threaded code written for the SPP-2000 will form the basis for our work on the Tera. Should interesting results be obtained, we will also investigate the porting of other codes of interest to NASA to the Tera.

GIANFRANCO CIARDO

Distributed State-Space Exploration Using a Shared-Memory Architecture

Practical continuous-time Markov models have a very large underlying state space, which, in general, must be explored and stored explicitly before attempting a numerical solution. In [1], we have studied approaches to perform this exploration in a message-passing environment, such as a network of workstations, in such a way that both the state space and the transition rate matrix of the Markov chain are partitioned over the available processors. The resulting speedup is quite good, and memory load balance issues have been addressed. However, the speedup of the numerical solution itself is not as good, due to the high latency of messages. Hence, we are now proposing to consider analogous distributed algorithms for a shared-memory environment.

In our message-passing approach, a partitioning function based on hashing or on information gathered from a short pre-exploration phase is used to determine the owner of a state. Then, each process can manage its portion of the state space independent of the others. In a shared memory environment, instead, a single data structure is used to store the entire state space and appropriate measures must be taken to avoid concurrent updates of key pieces of information. Of course, locks can be used to enforce mutual exclusion, but they incur a high overhead so their use should be limited. One fundamental observation that should enable us to achieve acceptable speedup is that, if the state space is organized as a search tree, most searches will not cause this tree to be modified. This is because, if the reachability graph of the model contains S states and T transitions between states, T searches will have to be performed, but only S of these will result in a new state. Since T is normally at least one order of magnitude larger than S , we will assume that "read access" to the search tree as the default, and only when an insertion is required, a lock will be issued, to ensure "write access". More efficiency might be achieved by adopting the multilevel data structure introduced in [2], since this would have the effect of limiting most of the search tree updates to lower levels, so that multiple concurrent read and write accesses could be performed, as long as they are non-overlapping.

This research has been done in collaboration with David Nicol (Dartmouth College) and is ongoing. A large implementation effort will be required before we can obtain meaningful timing results for the various approaches.

The Uniformization Algorithm in the Presence of Instantaneous States

The uniformization (or randomization) algorithm is perhaps the best-known approach for the transient analysis of continuous-time Markov chains (CTMCs). However, many formalisms, such as generalized-stochastic Petri nets (GSPNs), define an underlying semi-Markov process where the states are either timed, with an exponentially distributed sojourn time, or instantaneous, with a zero sojourn time. Only after

eliminating the instantaneous states a CTMC is obtained, which can then be studied using uniformization. We have recently introduced a class of solution approaches that avoid the storage of large matrices through the use of Kronecker operators, but, under certain conditions, their application forbids one to eliminate the instantaneous states, so only steady-state analysis could be performed in this case. In this research, we seek to extend the uniformization algorithm to a so-called “preservation” approach where the instantaneous states are kept throughout the computation.

The key idea of uniformization involves a transformation of the CTMC into a discrete-time Markov chain (DTMC) using a uniformizing rate q . This can still be performed on the type of semi-Markov processes we consider, but with a fundamental complication. In the original approach, we need to compute the probability that the DTMC performs k steps in the interval $(0, t)$, and this is simply the k -th entry of the Poisson distribution with parameter qt . In our case, instead, this holds only for steps in timed states, while an arbitrarily large number of steps can occur in the instantaneous states between any two visits to timed states. Unfortunately, there is no simple way to compute the probability of being in a times state i after k timed steps and an arbitrary number of instantaneous steps, for each i . A first approach to the solution of this problem is then to express these probabilities through an algorithm which requires k steps, each involving the solution of a linear system. We have been able to show that the complexity of each of these solutions is related to the length of instantaneous paths in the semi-Markov process, so that the proposed algorithm is acceptably fast if there are few short paths of this type, but it can be very slow in models where there are many cycles of instantaneous states.

We intend to proceed in two directions. First, we intend to investigate alternative algorithms that have a better complexity in the worst case. Then, we intend to implement our research results to gain timing results on practical models.

This work was done in collaboration with Kishor Trivedi (Duke University).

On the Use of Kronecker Operators for the Solution of Large Markov Models

The storage of the infinitesimal generator Q of a continuous-time Markov chain (CTMC) is the major obstacle to a numerical solution, since it requires unacceptably large amounts of memory for practical situations. Recently, several authors have proposed to store Q implicitly as a sum of Kronecker (or tensor) products of a set of much smaller matrices. However, previous proposals have severe limitations, both in the range of applicability and in the performance of the solution methods. We have greatly improved the state-of-the-art in both respects.

To extend the applicability of the method, we have moved away from the requirement of having a set of synchronized submodels, each of which must be itself an ergodic CTMC. Effectively, we have shown how to allow arbitrary Markovian structures to be synchronized in rather general ways. In addition, we also eliminated one of the earlier limitations, that synchronizations can only occur through timed events: instantaneous synchronization can now be managed through a discrete-time embedding. In the area of solution algorithms, we have earlier presented a set of vector-matrix multiplication algorithms, for matrices expressed as Kronecker products, and we have shown various tradeoffs when using these algorithms to solve a CTMC for steady-state, with a particular attention to methods that use the “actual state space” (instead of the much larger “potential state space” used by most previous authors). An unexpected result is that a multiplication by rows (as feasible in the Jacobi method) can achieve a much lower overhead than one by column (as required by the Gauss-Seidel or SOR method): hence we must choose between a method requiring, in general, more (but faster) iterations, or fewer (but slower) iterations.

A uniprocessor implementation of the above techniques is ongoing. A distributed implementation of a Kronecker-based approach is also highly appealing; one key advantage is that the matrix Q , being stored only implicitly, can be *de facto* duplicated on each processor at little additional cost.

TOM CROCKETT

Fast Algorithms for Color Image Quantization

Color quantization is the process of reducing the color resolution of a digital image, e.g., to compress the image or to display it on monitors with limited color capabilities. The goal is to find a limited color palette which will serve as an acceptable substitute for the much larger number of colors found in the original image.

In previous work with Shahid Bokhari and David Nicol, we modified the Parametric Binary Dissection algorithm for use in color image quantization, resulting in a new algorithm which we call Fast Adaptive Dissection, or FAD. The FAD algorithm produces good quality images at high speed. During the current reporting period, we summarized our results in an ICASE report and submitted them for publication. We also produced a companion video which illustrates the operation of the algorithm, including a computer-generated animation of the color space partitioning process.

In the future, we plan to adapt the sequential FAD algorithm for use in a parallel environment within the PGL rendering system. We are also interested in alternate algorithms which may present better opportunities for parallelization.

Parallel Polygon Rendering using Networks of Workstations

For the last several years, we have been developing a prototype parallel polygon rendering system known as PGL, supported under the auspices of NASA's HPCCP/CAS program. In February 1997, the first public release of PGL was made available to the HPCC community, with support for the IBM SP2 and Intel Paragon systems. During the same time period, clusters of networked workstations or PCs have grown in popularity as alternatives for certain classes of parallel applications. Consequently, there has been considerable interest from the user community in a version of PGL which would run in that environment.

In fulfillment of an HPCCP/CAS milestone, we have ported PGL to run on networks of Sun or SGI workstations, using any of several readily-available implementations of the MPI message passing standard. This provides the ability for distributed workstation-based computations to generate live visual output in exactly the same manner as applications which run on larger and more expensive parallel supercomputers. The primary difference is performance, which will generally be lower in the cluster environment, since communication overheads are higher and the number of available processors is usually modest.

To date, we have not had access to a workstation cluster environment which is suitable for conducting controlled performance experiments with PGL. The primary difficulty is finding a sufficiently large collection of systems with adequate networking infrastructure which can be scheduled in a dedicated mode. If the facilities become available, we plan to conduct experiments to characterize rendering rates and communication overheads, determine suitable buffering parameters, and examine scalability and parallel efficiency in the workstation environment.

We are also working to port PGL for use with newer versions of SGI's IRIX operating system, a necessary prerequisite for use with the Origin2000 systems which are expected to be installed as NASA's next CAS testbed.

Image Compression Schemes for Parallel Rendering Applications

While considerable effort has been directed at the development of efficient parallel rendering algorithms, less attention has been paid to the problem of delivering the resulting stream of images to remote users. The simple lossless compression techniques employed in the PGL rendering system have been adequate for use on local area networks, but they are too slow on congested long-haul routes across the Internet.

To address this problem, we have been developing more aggressive image compression strategies which satisfy certain criteria, including compatibility with parallel rendering algorithms, the exploitation of parallelism in the compression process, and rapid decompression at the receiving end. We are currently studying wavelet techniques, and have developed the outline of a parallel algorithm for wavelet compression in a distributed memory environment. Other compression methods under consideration include JPEG and FELICS.

Our immediate plans are to refine the details of the parallel wavelet algorithm, and to test it within the PGL system. We also plan a more detailed assessment of the JPEG and FELICS techniques for use in parallel rendering, and plan to develop an end-to-end performance model which will capture the important components in the image generation and delivery pipeline. The model will allow us to determine which techniques are most suitable in different scenarios of computational and networking performance.

This work was done in collaboration with Kunal Mukherjee and Amar Mukherjee of the University of Central Florida.

STEPHEN GUATTERY

Extending Graph Embedding Bounds to Incomplete Factor Preconditioners

The number of iterations required for convergence is an important measure of the performance of iterative methods such as conjugate gradient and preconditioned conjugate gradient. In most cases, this measure is difficult to determine; however, upper bounds can be determined from the spectral condition number, the ratio of the largest to smallest eigenvalues of the system. Calculating spectral condition numbers exactly can require substantial work and storage. Hence it is desirable to find a general method that can be applied to a wide range of matrices and preconditioners. In his Ph.D. thesis, Gremban used graph embeddings in bounding condition numbers for preconditioned systems where the matrix and preconditioner are symmetric, irreducibly diagonally dominant M-matrices (such matrices have positive diagonal and nonpositive off-diagonal entries). He also gave an extension that allows positive off-diagonal entries as long as the matrix remains diagonally dominant. His techniques are often easy to apply and often give good bounds. They can take advantage of the wide variety of embedding results developed in the study of networks. However, they apply only to a restricted set of matrices.

We have extended graph embedding techniques to symmetric positive definite systems where the preconditioner is not diagonally dominant. In particular, this allows us to bound the spectral condition number when the preconditioner is based on an incomplete Cholesky factorization. The extension is based on a splitting of the error matrix of the preconditioner and a careful application of the Support Lemma to the resulting matrices. In addition to providing useful bounds, this technique also provides intuitive insight about incomplete Cholesky and modified incomplete Cholesky preconditioners.

Several topics for further research remain. In particular, we hope to use these bounds as part of our work on computing good orderings for incomplete factorization. We also would like to use them in determining good heuristics for deciding which entries to use in an incomplete factorization.

This work was part of a joint effort with Alex Pothén of ICASE and ODU.

FABIO GUERINONI

Parallel PAB3D: Experiences Using MPI in a Large CFD Code

Parallel processing has been a trend in the aerospace industry for more than a decade. A number of systems have emerged which run on a number of processors. The objective of our particular problem was to parallelize a real application: The PAB3D code. This code incorporates multi-block/multi-zone techniques on a general purpose Reynolds-Averaged Navier-Stokes solver for propulsion and complex aerodynamic configurations and has several schemes for the RANS including turbulence models, gas mixtures, and multi-block capability.

The sequential code was initially developed at NASA in 1986 for supersonic jet exhaust flow analysis. We have had considerable success in designing a prototype message passing prototype for PAB3D using MPI (Message Passing Interface) which is becoming a standard. The development of the prototype took approximately ten weeks.

The parallelization of the code consisted of three phases. The first, perhaps the most laborious, consisted of the identification of critical sections of the code, i.e., those that cannot be executed in parallel. At the end of the phase, we ensured that each processor had access to the right data.

The second phase was the division of computation among the processors, prior to the actual computation. This involved the development of COMMSYS, MPI based set of routines to access the databases in PAB3D. By a combination non-blocking and blocking calls and the use of communication buffer, messages were exchanged among processors. Finally, in the third phase, processors independently computed conserved quantities and residuals. The turbulence model was disabled in the prototype, but the structure of the code will allow a similar treatment in the full code.

All computations were made on a mesh representing a nozzle involving nearly 1.3 three-dimensional nodes on two Unix clusters (file systems), involving 9 SGI workstations running IRIX 5.2. We used the LAM/MPI version 6.1 of the Ohio Supercomputer Center.

There are a number of implementation decisions that are still open. However, the experiences that we have gained so far, especially during the development of phase one, have led to a redesign of the code. Among the important issues we need to consider are: the best way to implement across-time iterations and many-to-one operations. The implementations of even simple instructions may differ greatly and may have an impact on the overall execution of the code.

More significant improvements should involve, a shrunk global block, automatic load balancing and distributed I/O.

MATTHEW HAINES

Software Support for Multi-Module Programming

Creating large programs that are composed from a set of previously-defined, smaller programs is desirable. For example, this technique is commonly used to create multidisciplinary applications and large-scale simulations. Unfortunately, from a software engineering perspective it is very difficult to combine such programs. Programmers wishing to do so must deal with different data formats, I/O specifications, and the problem of getting each independent module to act as a "server" for its functionality so that it can be externally controlled.

We are working to provide solutions to the software engineering problems present in creating multi-module programs. Our approach is to create a runtime library that provides all of the functionality needed to

i) turn ordinary programs into plug-and-play modules, and ii) organizing the modules into a single, coherent multi-module program. These operations are done at a high level of abstraction, and can reasonably be performed by a scientist or engineer. All of the lower-level communication issues are seamlessly handled by the runtime system. We have successfully built this runtime system and are in the process of creating several multi-module applications from real-world applications to test our approach.

We plan to continue testing our runtime system in its current form as well as using it to provide the underlying support for two other projects. One is the Opus project being developed at ICASE, the University of Vienna, and the University of Wyoming. The other is a java-based graphical tool for constructing multi-module programs that is being developed at ICASE.

This work was done in collaboration with Piyush Mehrotra, ICASE, and Hans Zima, University of Vienna.

VICTORIA INTERRANTE

Visualizing 3D Flow with Volume Line Integral Convolution

Efficient and effective techniques for visualizing numerical simulations of three-dimensional flow are necessary and useful for facilitating CFD research, and for aiding the communication of design concepts and consequences between researchers in different disciplines.

Based on insights from research in human visual perception, we have developed improved methods for conveying information about a 3D flow through a volume. We have shown how three-dimensional line integral convolution can be used to generate a solid texture of scan-converted streamlines, which can be rendered with visibility-impeding halos to emphasize depth discontinuities in the projection to any arbitrary 2D view, and can be colored to reflect the values of associated scalar quantities throughout the 3D region.

We would like to further explore the interaction of stereo and motion in facilitating the rapid comprehension of 3D flow information.

Visualizing Layered Surfaces in Volume Data

It can be difficult, in an ordinary rendering, to adequately perceive the 3D shape of a layered transparent surface, or to accurately judge its distance in depth from underlying opaque objects. Our objective is to facilitate the comprehensive understanding of multiple superimposed layers by enhancing the presentation of shape and depth information through the design and application of perceptually appropriate, sparse surface texture.

Based on insights from psychophysical research in shape and depth perception, we have developed a method for texturing arbitrary transparent surfaces with a set of evenly-distributed strokes that can intuitively convey relevant surface shape information. Our algorithm is fully automatic, and can be used to generate a single solid texture that applies to all isolevel surfaces in a smooth 3D distribution.

We would like to investigate the issues involved in using texture to portray more than two overlapping layers, and in particular to gain insight into the relative effectiveness of different texture parameters in facilitating the exclusive perceptual grouping of elements on one of two superimposed surfaces. We would also like to pursue a 2D, OpenGL implementation of our surface texturing algorithm.

DAVID KEYES

Parallelization of Implicit CFD Codes

The use of analysis codes in optimization often requires access to derivative information, e.g., the inverse action of the Jacobian of the state constraints. Since this implicitness is required in target use of the codes, it should be exploited in an implicit Newton-like solution process. Newton-Krylov-Schwarz (NKS) methods permit the data locality natively present in most PDE-based applications (CFD, CSM, CEM, etc.) to be exploited on distributed memory systems and deliver machine-zero satisfaction of the discrete conservation laws, together with the inverse action of the Jacobian at convergence. Our goal has therefore been to demonstrate NKS methods on realistic 3D NASA flow configurations, and to gain experience tuning their many parameters so that analysts and optimizers are attracted to their use.

We work directly with legacy F77 sources for definition of the problem discretization, and we explicitly manage domain partitioning, construct domain-decomposition (overlapping Schwarz) preconditioners, and tune the successively nested algorithmic levels: (1) timestepping (with CFL approaching infinity for steady-state solutions), (2) Newton convergence on each time step, (3) Krylov convergence on each Newton step, and (4) preconditioner strength. Parallel NKS-based ports are implemented with the MPI-based Portable, Extensible Toolkit for Scientific Computing (PETSc). We have completed ports to distributed memory (SP and T3E) and distributed shared memory (Origin) of structured-grid and unstructured-grid Euler codes, and tested both on an ONERA M6 wing geometry. Performance has been tuned for both the memory hierarchy and the communication network. For example, parallel sustained performance of approximately 70 Mflop/s per processor and relative parallel efficiency of 75 per cent from 8 to 80 processors have been attained on the IBM SP for a 1.4 million degree-of-freedom problem. Parallel sustained performance of approximately 65 Mflop/s per processor and relative parallel efficiency of 83 per cent from 128 to 512 processors have been attained on the Cray T3E for an 11 million degree-of-freedom problem. The per processor performance of the unstructured-grid code is approximately eight times faster than the original vector-oriented version, running sequentially, and is comparable to the same algorithm applied to a structured-grid Euler code running in parallel. (Cache-oriented restructurings to achieve this level of performance included subdomain blocking, point blocking, field interleaving, and edge reordering.) The high parallel efficiency is possible because iteration count degrades only weakly as the domain is partitioned for concurrent preconditioner application.

While additional physical fidelity is inserted into the parallel analysis codes, we will work with their user communities to make workstation network and HPCCP testbed use "friendly." This will include attention to parallel I/O, checkpointing, and mesh sequencing. We will work on improved partitioning methods, in order to load balance the sum of local work and local communication, rather than just the local work, as at present. We will also work on partitioning methods that preserve strong flow-oriented couplings, for improved convergence rates. Finally, we will explore the additional task parallelism available in use of the parallel analysis codes in optimization.

This work is joint with W. Kyle Anderson of NASA, Dinesh Kaushik and Nilan Karunaratne of ODU, and William D. Gropp, Lois C. McInnes, and Barry F. Smith of Argonne National Laboratory.

KWAN-LIU MA

Java-Based User Interface Design for Remote and Collaborative Volume Visualization

Direct volume rendering is the display of data sampled in three dimensions and can generate more realistic and informative visualization of the data than conventional visualization techniques like surface

rendering and contour plots. While previous research efforts in volume rendering have focused mostly on the development and optimization of rendering algorithms, appropriate user interface design is needed now to make volume rendering a truly practically useful visualization tool. The goal of this research is to develop an intuitive user interface which makes the process of volumetric data exploration as efficient and useful as possible.

Our design takes advantage of the current web technology and has been implemented in Java to support remote as well collaborative visualization. The design consists of three parts: graphical user interface (gui), visualization process and remote/collaborative visualization. The main gui components include interfaces for displaying images, calculating statistics of the data, defining the color and opacity transfer functions, and tracking visualization steps. The visualization process keeps track of all the rendering and viewing parameters of each visualization made using a directed graph. This graph can be displayed and its components may be selected to perform further visualization calculations. Remote/collaboration visualization is supported by a data downloading mechanism, an annotation capability, the gui and the tracking ability. A prototype system has been implemented and is used as a testbed to explore a variety of user interface approaches for remote/collaborative volume visualization. The remote functionality makes volume visualization possible even on machines with a small amount of memory, processing power or disk space. The collaborative capability allows scientists to record and communicate their ideas and findings in a more efficient manner.

We are currently implementing an automatic animation capability based on user's previous visualization process. We plan to perform user study to improve the feel and look of the system. We will investigate incremental loading and rendering of the data for remote visualization. We will look into the possibility of distributing rendering tasks to multiple workstations. We will also develop additional features for supporting collaborative visualization.

This work was done in collaboration with James Patten of University of Virginia.

PIYUSH MEHROTRA

Multithreaded System for Distributed Environments

Traditionally, lightweight threads are supported only within the single address space of a process, or in shared memory environments with multiple processes. Likewise, interprocess communication systems do not currently allow messages to be sent directly entities within a process. The objective of this project is build a system which combines standard interfaces for lightweight threads, pthreads, and interprocess communication, MPI, to support point-to-point communication between any two threads in a distributed memory system.

The Chant runtime system has been built using layers: point-to-point communication, remote service requests, remote thread operations. In the last year we have added a layer to support load balancing via migration of threads. In contrast to other thread migration systems, we provide migration in the presence of pointers, i.e., along with the stack the thread heap is also migrated. This allows the pointers to point to valid data even after the migration. The load balancing layer provides facilities for computing the current workloads of the processors, figuring which threads should be moved and to what processors. The system also keeps track of the amount of communication being generated by each thread so that it can be used in the decision making process. The underlying thread migration layer takes care of the actual motion of the threads. The system is designed such that default routines can be replaced by the user with specialized routines for carrying out the various functions.

We are currently debugging and optimizing the load balancing layer. We are currently implementing several thread-based parallel codes and will use these codes to drive the optimizations.

This work was done in collaboration with David Cronk (ICASE) and Matthew Haines (University of Wyoming).

Evaluation and Extension of High Performance Fortran

The stated goal of High Performance Fortran (HPF) was to “address the problems of writing data parallel programs where the distribution of data affects performance.” We have been using data parallel codes of interest to NASA to evaluate the effectiveness of the features of HPF and to suggest extensions for the language.

We have continued our exploration of multi-block CFD codes using a model version of the TLNS3d-MB code developed at NASA Langley. We have examined three different distributions of the data each exploiting different levels of parallelism in the code being able to compile two of these versions using currently available HPF compilers. The speedup of the code generated by the HPF compilers compares favorably with speedup obtained by the hand coded MPI versions of the model code. However, the single node performance of the HPF generated code is below par.

We have also been evaluating the use of HPF for unstructured grid codes which uses indirection arrays to access data. HPF compilers use a runtime inspector-executor strategy to generate the required communication schedules. We have devised new directives which allow the user to control when to generate and reuse such schedules. Preliminary performance figures, based on hand-coded translations, indicate that the performance of using this approach can be substantial.

We have also used a classic nonlinear elliptic PDE, the Bratu problem, to compare the performance of HPF with that obtained by using the PETSc libraries. The test case is solved with a Newton-Krylov-Schwarz iterative method, with a second-order finite difference discretization, and explicit computation of the Jacobian. Both approaches require the explicit specification of data partitioning. With relatively modest effort, we obtain similar and reasonable performance using both paradigms.

A new version of HPF (HPF 2.0) which includes several extensions resulting from our experiments, has been released. We plan to continue the evaluation of the new features of HPF as they are incorporated in the compilers.

This work was done in collaboration with M. Hayder (ICASE), D. Keyes (ODU and ICASE), K. Roe (ICASE), J. Van Rosendale (National Science Foundation and ICASE), and H. Zima (University of Austria).

Arcade: A Distributed Computing Environment for ICASE

Distributed heterogeneous computing is being increasingly applied to a variety of large size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consists of multiple heterogeneous modules interacting with each other to solve the problem at hand. Such applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI based environment which supports the multi-user design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines and parallel architectures.

The overall goal is to design an interface which is easy to use, easily accessible, and portable. We are planning to leverage off of technologies where available to achieve these goals. In particular, integrating the interface in a web browser, e.g., Netscape, provides users with a familiar interface on desktops ranging

from Unix based workstations, to Windows based PCs, and Macintoshes. The Aracde system consists of the following sub-environments: *application design interface*, *resource allocation and execution interface* and *monitoring and steering interface*. The architecture is a three-tier architecture. The front-end, based on a Java, provides the user interface. The middle tier consists of the user interface server which provides the logic to process the user input and to interact with application modules running on a heterogeneous set of machines. Application execution is monitored by the Execution Controller which interacts Process Controllers on the individual resources. The PCs along with application modules form the last tier of the architecture.

The current prototype implementation allows single users to specify non-hierarchical applications through a textual or a simple visual system. The resources have to be statically specified by the user. The system supports the execution of the modules on a network of workstations in a single domain. The execution status of the application can be monitored by multiple users simultaneously.

We are currently expanding the system to incorporate all the facilities envisioned in our architecture in a phased approach, an approach where the driving force is the user of the system. In particular, we are currently working a visual interface for specification of the applications, mechanisms to integrate other approaches to specifying multidisciplinary applications, such as Opus, and also on automatic scheduling of the application modules.

This work was done in collaboratin with Zhikai Chen, K. Maly, P.K. Vangala, and M. Zubair (all affiliated with ODU and ICASE).

WILLIAM MITCHELL

Parallel Adaptive Multigrid

Adaptive multilevel methods combine multigrid iteration with adaptive refinement to produce fast $O(N)$ solutions of differential equations that exhibit locally strong variations. These methods have been shown to be efficient on sequential computers, however many problems of interest require parallel computers. While adaptive refinement, multigrid and parallel computers are individually achieving wide spread use in applications, the combination of all three is still a research topic due to the inherent irregularities of adaptive grids and multigrid, and the requirement of frequent communication. The full domain partition (FuDoP) was recently developed to reduce the communication in adaptive multilevel methods, and has been implemented in the Parallel Hierarchical Adaptive MultiLevel (PHAML) code for experimentation in the context of 2D linear elliptically dominated systems. Optimal convergence of the FuDoP multigrid method has been observed numerically. Current research is investigating the theoretical rate of convergence.

The FuDoP multigrid method can be expressed in the Schwarz framework of domain decomposition methods. During my two weeks at ICASE, theoretical analysis of the convergence rate of the FuDoP multigrid method was examined. The method can be interpreted as an additive Schwarz method (ASM) applied to the overlapping subdomains, with a multigrid subdomain solver. Under the assumption that the abstract Schwarz framework applies, a convergence rate for the FuDoP multigrid solver was established in terms of the constants from the abstract framework and the convergence rate of the sequential multigrid method. This required a minor modification of the method to symmetrize the subdomain solver. While searching for bounds on the Schwarz constants, it was discovered that this is not sufficient for the method to fit under the given abstract Schwarz framework, due to in the inherent nonsymmetry of the restriction/prolongation operator pair in the additive Schwarz portion of the method. However, this nonsymmetry supports a connection between this method and the new Restricted Additive Schwarz method (RASM) of X.-C. Cai, which

is experimentally known to be superior to symmetric ASM in terms of both communication frequency and convergence rate. Theoretical analysis of RASM is currently in progress. The understanding of the relationship between these methods should allow us to apply the theoretical analysis of RASM to the FuDoP multigrid method.

As theoretical developments in RASM become available, they will be studied for applicability to the FuDoP multigrid method. The PHAML code will be extended to broader classes of differential equations (nonlinear, time dependent, 3D, systems), and demonstrated with practical application problems.

CAN OZTURAN

Parallel Mesh Environment (PME)

Parallel adaptive PDE solvers on two- and three-dimensional complex geometry need a greater repertoire of mesh entity adjacency relationships for h-refinement, parallel refinement, coarsening and dynamic load balancing routines. The implementation of these routines need specialized data structures which provide fast updates not only on a single processor but also during entity migrations among processors. The objective of this project is to develop an environment, Parallel Mesh Environment (PME), for distributed unstructured grid manipulations.

During the past year, the following goals were accomplished:

- Developed a data parallel adaptive mesh refinement algorithm based on longest edge bisection. Our algorithm has logarithmic worst case complexity and improves the earlier algorithms which had linear worst case complexity.
- Parallelized a NASA Langley Aeroacoustic code (written in C++) by using PME. This project was carried out in collaboration with Dr. Harold Atkins from NASA Langley and Abdelkader Baggag (graduate fellow from Univ. of Minnesota).
- Interfaced PME with the METIS mesh partitioner software. This interface enables calling of METIS directly from PME and a follow-up migration of mesh data structures based on the output of METIS software.
- Coded a bisection-based tetrahedral refinement algorithm for parallel implementation. This code is currently in debugging phase.
- Implemented examples to show how PME and PETSc can be coupled to solve partial differential equations on unstructured grids.

Short term plans include documentation and release of PME and coding of adaptive mesh refinement examples with PME and PETSc.

DEBORAH PILKEY

Computation of a Damping Matrix for Finite Element Model Updating

Damping matrix identification is an important tool for diagnostics and health monitoring of structures. The energy dissipation typically provided by a damaged structure is best characterized by its damping matrix. Research into the identification of damping matrices is in its infancy because of the physical, numerical and computational challenges faced with such a problem. The objective of this research is to investigate the computational issues faced with damping matrix identification including model expansion and updating, as well as the inverse eigenvalue problem, itself.

Investigations into the computational issues of model updating are under way. High Performance Fortran is being used to show that large finite element models can be verified with data from smaller experimental models. Several model reduction and expansion methods have been coded and are running to determine the most appropriate for use with high performance computing. In addition, issues of damping matrix identification and the associated inverse eigenvalues problems have been explored using High Performance Fortran. Benefits of iterative algorithms over direct methods have been realized. The work thus far has shown that high performance computing can be a great benefit in this area of structural mechanics.

Model updating algorithms are being sought that take damping into consideration. These will be investigated for not only engineering accuracy, but also efficiency in large-scale applications.

ALEX POTHEN

Ordering Irregular Data Structures to Reduce Cache Misses

Computations on unstructured meshes and irregular, sparse matrices or graphs require greater running times than computations on structured meshes and dense matrices with comparable operation counts. The primary reason for this is that the irregular data structures lead to many more cache misses on modern superscalar architectures with multiple issue, branch prediction, and replicated functional units. We consider the creation of algorithms and software to order the computations on the irregular data structures to reduce cache misses and thereby improve performance of the computations.

We have designed two algorithms for computing "locality-inducing" orderings on irregular graphs and meshes. The first algorithm is an improved implementation of a combinatorial algorithm due to Sloan. We described an algorithm whose running time is bounded by the number of edges (nonzeros) in the sparse graph (matrix). Consequently the new implementation is an order of magnitude faster than the earlier implementations. We also improved the quality of the algorithm by identifying two classes of problems and choosing parameters in the algorithm appropriately for each class. This algorithm currently provides good quality with low execution times, and is the best combinatorial algorithm for the problem at this time. We have also improved an algebraic (spectral) algorithm that we had designed in earlier work by using a post-processing phase involving the Sloan algorithm. The hybrid algorithm further reduces a metric called the wavefront over the spectral and Sloan algorithms but requires a greater running time. Examples show that these improved orderings lead to frontal solvers that are faster by an order of magnitude on advanced architectures. An important application of this work is in computing incomplete factorization preconditioners for iterative solvers for partial differential equations. The object-oriented software that we have developed is available with three interfaces: Matlab (a commercial programming environment and language for scientific computing and visualization), PETSc (Portable, Extensible Toolkit for Scientific Computing), and as a stand-alone code. It has been distributed to users at NASA Ames and to several academic researchers.

We are continuing our study by measuring the cache misses from these orderings for unstructured mesh computations, and how different cache organizations influence the desirable orderings.

This is joint work with Gary Kumfert, current Ph.D. student at Old Dominion University, who is also affiliated with ICASE.

Parallel Algorithms for Incomplete Factorization Preconditioners

The parallel computation of robust preconditioners is a priority for solution of large systems of equations in unstructured grids and other applications. We are developing new algorithms and software that can

compute incomplete factorization preconditioners in time proportional to the number of floating point operations and memory accesses. These algorithms are also parallelizable, and we will consider implementations on parallel machines.

We have developed a structure theory based on paths in the adjacency graph of the matrix to predict where zero elements become nonzeros in incomplete factorization (fill elements). A level function is used in incomplete factorization to control the number of fill elements, and we relate the level of fill to lengths of appropriately defined paths in the adjacency graph. This result permits us to search in the neighborhood of a vertex in the graph to predict all fill elements associated with that vertex. We have designed three variants of these algorithms and have proved that they have a smaller running time complexity than currently used algorithms for computing incomplete factorizations. The more efficient algorithms make use of the concept of transitive reduction of directed graphs (symmetric problems) and symmetric reduction of directed graphs (unsymmetric problems) in order to search for paths in smaller graphs. Our current object-oriented implementations show that the new algorithms are faster than the currently used implementations.

We will continue to develop algorithms and software for fast computation of preconditioners, and investigate parallel implementations. The parallel algorithms will make it possible to compute high level incomplete factorization preconditioners for parallel Krylov-space-based iterative solvers.

This is joint work with David Hysom, current Ph.D. student at Old Dominion University, who is also affiliated with ICASE.

A Parallel Sparse Indefinite Solver

Large, sparse, symmetric indefinite systems of equations occur in computational structural mechanics, electromagnetics, and linear and nonlinear programming. Our goal is to develop parallel algorithms and software for solving sparse, symmetric indefinite linear equations on serial machines as well as distributed-memory multiprocessors.

We continued developing the first parallel solver known to us for sparse indefinite systems of equations. The multifrontal method was used to organize the factorization to make effective use of cache and memory accesses. New parallel algorithms and dynamic data structures were developed to deal with the irregular computation caused by sparsity and numerical pivoting. MPI was employed for portability. We have real and complex versions of the code for our serial implementation. We have used this code to solve problems from structural analysis, fluid flow, Helmholtz problems in computational acoustics and computational electromagnetics, and linear programs. On all these problems, the solver is able to compute accurate solutions; on many of these problems, a positive definite solver would fail, or would require a much greater running time because of greater fill.

We are tuning our solver for improving the performance, and implementing ordering algorithms suitable for indefinite problems. We are collaborating with experts in acoustics and electromagnetics to solve indefinite problems from these application areas.

This is joint work with Florin Dobrian of Old Dominion University.

HANS ZIMA

High Performance Fortran for Irregular Data Parallel Problems

High Performance Fortran (HPF) provides the user with a high-level language interface for programming scalable parallel architectures and delegates to the compiler the task of producing an explicitly parallel

message-passing program. The current language standard (HPF-2) does not support some functionality that we consider essential for efficiently handling irregular applications. In previous work performed jointly at ICASE and the University of Vienna, we addressed this problem by defining language extensions for the explicit control of communication schedules. During this reporting period, we have generalized this set of extensions to provide support for irregular halos. Furthermore, explicit schedule control was implemented in the framework of the Vienna Fortran Compiler (VFC) and successfully applied to a set of benchmark codes on a Meiko CS-2 platform.

Future work will evaluate the present language extensions and study optimizations of the VFC runtime system supporting these features.

This work was performed with Siegfried Benkner and Viera Sipkova at the University of Vienna.

Integration of Task and Data Parallelism Using OPUS

Multidisciplinary applications (MDAs) require task parallelism in addition to the data parallelism offered by languages such as HPF. In previous work, ICASE and the University of Vienna have jointly developed a language called OPUS, which provides an object-based layer for the coordination of multiple HPF programs operating in a parallel or heterogeneous environment. During this reporting period, we have refined the language definition, produced a Reference Manual, designed an implementation and started work on the compilation and runtime system for the language. Furthermore, we have studied the relationship of OPUS to other approaches addressing this problem such as Linda, PCF, Orca, and CC++.

Future work will produce a complete implementation of OPUS, in parallel to an evaluation of the language guided by the results of its application and the experiences gained in the implementation effort. We will study optimizations, in particular with respect to parallel method executions and special implementation strategies for a class of "light" SDAs, which do not require the entire functionality of OPUS and HPF. In addition, we plan to integrate OPUS with the Arcade system developed at ICASE and consider the potential role of the language in a future graphical design environment for MDAs.

This work was performed together with Piyush Mehrotra and John Van Rosendale at ICASE, Matt Haines at the University of Wyoming, and Erwin Laure at the University of Vienna.

REPORTS AND ABSTRACTS

Wilson, Robert V., and Ayodeji O. Demuren: *Numerical simulation of turbulent jets with rectangular cross-section.* ICASE Report No. 97-1, (NASA CR-201642), January 31, 1997, 19 pages. To be submitted to the ASME Journal of Fluids Engineering.

Three-dimensional turbulent jets with rectangular cross-section are simulated with a finite-difference numerical method. The full Navier-Stokes equations are solved at low Reynolds numbers, whereas at the high Reynolds numbers filtered forms of the equations are solved along with a subgrid scale model to approximate effects of the unresolved scales. A 2-N storage, third-order Runge-Kutta scheme is used for temporal discretization and a fourth-order compact scheme is used for spatial discretization. Computations are performed for different inlet conditions which represent different types of jet forcing. The phenomenon of axis-switching is observed, and it is confirmed that this is based on self-induction of the vorticity field. Budgets of the mean streamwise velocity show that convection is balanced by gradients of the Reynolds stresses and the pressure.

Somani, Arun K.: *Reliability modeling of structured systems: Exploring symmetry in state-space generation.* ICASE Report No. 97-2, (NASA CR-201643), January 31, 1997, 30 pages. To be submitted to Sigmetrics.

A large number of systems are implemented using regular interconnected topologies. Markov analysis of such systems results in large state spaces. We explore symmetry, in particular rotational and permutational, of such systems to achieve a significant reduction in the size of the state space required to analyze them. The resulting much smaller state spaces allow analyses of very large systems. We define equivalent classes of states and develop an algorithm to generate small state spaces and the corresponding Markov chain for systems with permutation symmetries. The state space generation process is also simplified. We demonstrate our technique using several examples. Our technique is very useful in the exact analysis of large systems.

Somani, Arun K., and Kishor S. Trivedi: *Boolean algebraic methods for phased-mission system analysis.* ICASE Report No. 97-3, (NASA CR-201644), March 4, 1997, 21 pages. Submitted to IEEE Transactions on Reliability.

Most reliability analysis techniques and tools assume that a system is used for a mission consisting of a single phase. However, multiple phases are natural in many missions. The failure rates of components, system configuration, and success (failure) criteria may vary from phase to phase. In addition, the duration of a phase may be deterministic or random. We describe a new technique for phased-mission system reliability analysis based on Boolean algebraic methods. Our technique is computationally efficient and is applicable to a large class of systems for which the failure criterion in each phase can be expressed as a fault tree (or an equivalent representation). Our technique avoids state space explosion that commonly plague Markov chain-based analysis. We develop a phase algebra to account for the effects of variable configurations and failure criteria from phase to phase. Our technique yields exact (as opposed to approximate) results. We demonstrate the use of our technique by means of an example and present numerical results to show the effects of mission phases on the system reliability.

Haines, Matthew: *On designing lightweight threads for substrate software*. ICASE Report No. 97-4, (NASA CR-201645), March 4, 1997, 20 pages. To appear in the Proceedings of USENIX Technical Conference, Anaheim, CA, January 1997.

Existing user-level thread packages employ a “black box” design approach, where the implementation of the threads is hidden from the user. While this approach is often sufficient for application-level programmers, it hides critical design decisions that system-level programmers must be able to change in order to provide efficient service for high-level systems. By applying the principles of Open Implementation Analysis and Design, we construct a new user-level threads package that supports common thread abstractions and a well-defined meta-interface for altering the behavior of these abstractions. As a result, system-level programmers will have the advantages of using high-level thread abstractions without having to sacrifice performance, flexibility, or portability.

Ciardo, Gianfranco, and Andrew S. Miner: *Storage alternatives for large structured state spaces*. ICASE Report No. 97-5, (NASA CR-201646), March 4, 1997, 22 pages. Submitted to Performance Tools '97.

We consider the problem of storing and searching a large state space obtained from a high-level model such as a queuing network or a Petri net. After reviewing the traditional technique based on a single search tree, we demonstrate how an approach based on multiple levels of search trees offers advantages in both memory and execution complexity, and how solution algorithms based on Kronecker operators greatly benefit from these results. Further execution time improvements are obtained by exploiting the concept of “event locality.” We apply our technique to three large parametric models, and give detailed experimental results.

Aurovillian, Alok, Hong Zhang, and Malgorzata M. Wiecek: *A bookkeeping strategy for multiple objective linear programs*. ICASE Report No. 97-6, (NASA CR-201647), March 14, 1997, 14 pages. Submitted to the Journal of Computers and Operations Research.

This paper discusses the bookkeeping strategies for solving large multiple objective linear programs (MOLPs) on ADBASE, a well developed sequential software package, and on a parallel ADBASE algorithm. Three representative list creation schemes were first analyzed and tested. The best of them, Binary Search with Insertion Sort (BSIS), was selected to be incorporated into ADBASE and the parallel ADBASE algorithm. The resulting new bookkeeping strategy was then tested in ADBASE as well as implemented in the parallel ADBASE algorithm. The parallel implementations were carried out on an Intel Paragon multiprocessor. Computational results show that the new bookkeeping strategy for maintaining a list of efficient solutions significantly speeds up the process of solving MOLPs, especially on parallel computers.

Rubinstein, Robert, and Ye Zhou: *Time correlations and the frequency spectrum of sound radiated by turbulent flows*. ICASE Report No. 97-7, (NASA CR-201648), March 14, 1997, 24 pages. Submitted to Physics of Fluids.

Theories of turbulent time correlations are applied to compute frequency spectra of sound radiated by isotropic turbulence and by turbulent shear flows. The hypothesis that Eulerian time correlations are dominated by the sweeping action of the most energetic scales implies that the frequency spectrum of the sound radiated by isotropic turbulence scales as ϕ^4 for low frequencies and as $\phi^{-4/3}$ for high frequencies. The

sweeping hypothesis is applied to an approximate theory of jet noise. The high frequency noise again scales as $\sigma^{-4/3}$, but the low frequency spectrum scales as σ^2 . In comparison, a classical theory of jet noise based on dimensional analysis gives σ^{-2} and σ^2 scaling for these frequency ranges. It is shown that the σ^{-2} scaling is obtained by simplifying the description of turbulent time correlations. An approximate theory of the effect of shear on turbulent time correlations is developed and applied to the frequency spectrum of sound radiated by shear turbulence. The predicted steepening of the shear dominated spectrum appears to be consistent with jet noise measurements.

Somani, Arun K., and Allen M. Sansano: *Minimizing overhead in parallel algorithms through overlapping communication/computation*. ICASE Report No. 97-8, (NASA CR-201649), March 14, 1997, 28 pages. Submitted to IEEE Transactions on Parallel and Distributed Computing.

One of the major goals in the design of parallel processing machines and algorithms is to reduce the effects of the overhead introduced when a given problem is parallelized. A key contributor to overhead is communication time. Many architectures try to reduce this overhead by minimizing the actual time for communication, including latency and bandwidth. Another approach is to hide communication by overlapping it with computation. This paper presents the Proteus parallel computer and its effective use of communication hiding through overlapping communication/computation techniques. These techniques are easily extended for use in compiler support of parallel programming. We also address the complexity or rather simplicity, in achieving complete exchange on the Proteus Machine.

Anderson, W. Kyle, and V. Venkatakrishnan: *Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation*. ICASE Report No. 97-9, (NASA CR-201650), March 14, 1997, 46 pages. Presented at the 35th AIAA Aerospace Science Meeting, January 6-10, 1997, Reno, NV; also submitted to Computers in Fluids.

A continuous adjoint approach for obtaining sensitivity derivatives on unstructured grids is developed and analyzed. The derivation of the costate equations is presented, and a second-order accurate discretization method is described. The relationship between the continuous formulation and a discrete formulation is explored for inviscid, as well as for viscous flow. Several limitations in a strict adherence to the continuous approach are uncovered, and an approach that circumvents these difficulties is presented. The issue of grid sensitivities, which do not arise naturally in the continuous formulation, is investigated and is observed to be of importance when dealing with geometric singularities. A method is described for modifying inviscid and viscous meshes during the design cycle to accommodate changes in the surface shape. The accuracy of the sensitivity derivatives is established by comparing with finite-difference gradients and several design examples are presented.

Burns, John A., and Belinda B. King: *A note on the mathematical modelling of damped second order systems*. ICASE Report No. 97-10, (NASA CR-201657), March 14, 1997, 12 pages. Submitted to the Journal of Mathematical Systems, Estimation and Control.

This note is concerned with the formulation of a damped second order system as a first order dynamical system on a product space. This problem comes from the desire to have explicit representations of the infinitesimal generator of the first order system and, in particular, of the domain of this operator. This analysis is motivated by the need to find specific representations for Riccati operators that can be used in

the development of computational schemes for hyperbolic control problems. The approach we take here is based on a natural factorization of the differential operators that define the second order model.

Sun, Xian-He, and Yu Zhuang: *A high-order direct solver for Helmholtz equations with Neumann boundary conditions*. ICASE Report No. 97-11, (NASA CR-201658), March 26, 1997, 27 pages. Submitted to the International Conference on Supercomputing.

In this study, a compact finite-difference discretization is first developed for Helmholtz equations on rectangular domains. Special treatments, then, are introduced for Neumann and Neumann-Dirichlet boundary conditions to achieve accuracy and separability. Finally, a Fast Fourier Transform (FFT) based technique is used to yield a fast direct solver. Analytical and experimental results show this newly proposed solver is comparable to the conventional second-order elliptic solver when accuracy is not a primary concern and is significantly faster than that of the conventional solver if a highly accurate solution is required. In addition, this newly proposed fourth order Helmholtz solver is parallel in nature. It is readily available for parallel and distributed computers. The compact scheme introduced in this study is likely extendible for sixth-order accurate algorithms and for more general elliptic equations.

Lewis, Robert Michael: *A nonlinear programming perspective on sensitivity calculations for systems governed by state equations*. ICASE Report No. 97-12, (NASA CR-201659), March 26, 1997, 37 pages. Submitted to SIAM Review.

This paper discusses the calculation of sensitivities, or derivatives, for optimization problems involving systems governed by differential equations and other state relations. The subject is examined from the point of view of nonlinear programming, beginning with the analytical structure of the first and second derivatives associated with such problems and the relation of these derivatives to implicit differentiation and equality constrained optimization. We also outline an error analysis of the analytical formulae and compare the results with similar results for finite-difference estimates of derivatives. We then attend to an investigation of the nature of the adjoint method and the adjoint equations and their relation to directions of steepest descent. We illustrate the points discussed with an optimization problem in which the variables are the coefficients in a differential operator.

Chang, H.-C., D. Gottlieb, M. Marion, and B.W. Sheldon: *Mathematical analysis and optimization of infiltration processes*. ICASE Report No. 97-13, (NASA CR-201660), March 26, 1997, 19 pages. Submitted to the Journal of Computational Physics.

A variety of infiltration techniques can be used to fabricate solid materials, particularly composites. In general these processes can be described with at least one time dependent partial differential equation describing the evolution of the solid phase, coupled to one or more partial differential equations describing mass transport through a porous structure. This paper presents a detailed mathematical analysis of a relatively simple set of equations which is used to describe chemical vapor infiltration. The results demonstrate that the process is controlled by only two parameters, α and β . The optimization problem associated with minimizing the infiltration time is also considered. Allowing α and β to vary with time leads to significant reductions in the infiltration time, compared with the conventional case where α and β are treated as constants.

Leutenegger, Scott T., Jeffrey M. Edgington, and Mario A. Lopez: *STR: A simple and efficient algorithm for R-tree packing*. ICASE Report No. 97-14, (NASA CR-201661), March 31, 1997, 31 pages. To appear in the 1997 International Conference on Data Engineering.

In this paper, we present the results from an extensive comparison study of three R-tree packing algorithms, including a new easy to implement algorithm. The algorithms are evaluated using both synthetic and actual data from various application domains including VLSI design, GIS (tiger), and computational fluid dynamics. Our studies also consider the impact that various degrees of buffering have on query performance. Experimental results indicate that none of the algorithms is best for all types of data. In general, our new algorithm requires up to 50% less space than previously proposed algorithm for point and region queries on uniformly distributed or mildly skewed point and region data, and approximately the same for highly skewed point and region data.

Jones, Jim E., and N. Duane Melson: *A note on multi-block relaxation schemes for multigrid solvers*. ICASE Report No. 97-15, (NASA CR-201662), March 31, 1997, 12 pages. To be submitted to the 8th Copper Mountain Conference on Multigrid Methods.

Efficient and robust multigrid solvers for anisotropic problems typically use either semi-coarsened grids or implicit smoothers - line relaxation in 2D and plane relaxation in 3D. However, both of these may be difficult to implement in codes using multi-block structured grids where there may be no natural definition of a global 'line' or 'plane'. These multi-block structured grids are often used in fluid dynamic applications to capture complex geometries and/or to facilitate parallel processing. In this paper, we investigate the performance of multigrid algorithms using implicit smoothers within the blocks of a such a grid. By looking at a model problem, the 2-D anisotropic diffusion equation, we show that true multigrid efficiency is achieved only when the block sizes are proportional to the strength of the anisotropy. Further, the blocks must overlap and the size of the overlap must again be proportional to the strength of the anisotropy.

Jones, J.E., Z. Cai, S.F. McCormick, and T.F. Russell: *Control-volume mixed finite element methods*. ICASE Report No. 97-16, (NASA CR-201663), March 31, 1997, 28 pages. Submitted to Computational Geosciences.

A key ingredient in simulation of flow in porous media is accurate determination of the velocities that drive the flow. Large-scale irregularities of the geology (faults, fractures, and layers) suggest the use of irregular grids in simulation. This paper presents a control-volume mixed finite element method that provides a simple, systematic, easily implemented procedure for obtaining accurate velocity approximations on irregular block-centered grids. The control-volume formulation of Darcy's law can be viewed as a discretization into element-sized "tanks" with imposed pressures at the ends, giving a local discrete Darcy law analogous to the block-by-block conservation in the usual mixed discretization of the mass-conservation equation. Numerical results in two dimensions show second-order convergence in the velocity, even with discontinuous anisotropic permeability on an irregular grid. The method extends readily to three dimensions.

Horton, Graham: *On the multilevel solution algorithm for Markov chains*. ICASE Report No. 97-17, (NASA CR-201671), March 31, 1997, 24 pages. Presented at the 1996 Copper Mountain Conference on Iterative Methods; submitted to SISC.

We discuss the recently introduced multilevel algorithm for the steady-state solution of Markov chains. The method is based on an aggregation principle which is well established in the literature and features a multiplicative coarse-level correction. Recursive application of the aggregation principle, which uses an operator-dependent coarsening, yields a multi-level method which has been shown experimentally to give results significantly faster than the typical methods currently in use. When cast as a multigrid-like method, the algorithm is seen to be a Galerkin-Full Approximation Scheme with a solution-dependent prolongation operator. Special properties of this prolongation lead to the cancellation of the computationally intensive terms of the coarse-level equations.

Babin, A., A. Mahalov, B. Nicolaenko, and Y. Zhou: *On the asymptotic regimes and the strongly stratified limit of rotating Boussinesq equations*. ICASE Report No. 97-18, (NASA CR-201672), March 31, 1997, 43 pages. Submitted to Theoretical and Computational Fluid Dynamics.

Asymptotic regimes of geophysical dynamics are described for different Burger number limits. Rotating Boussinesq equations are analyzed in the asymptotic limit of strong stratification in the Burger number of order one situation as well as in the asymptotic regime of strong stratification and weak rotation. It is shown that in both regimes horizontally averaged buoyancy variable is an adiabatic invariant for the full Boussinesq system. Spectral phase shift corrections to the buoyancy time scale associated with vertical shearing of this invariant are deduced. Statistical dephasing effects induced by turbulent processes on inertial-gravity waves are evidenced. The 'split' of the energy transfer of the vortical and the wave components is established in the Craya-Herring cyclic basis. As the Burger number increases from zero to infinity, we demonstrate gradual unfreezing of energy cascades for ageostrophic dynamics. The energy spectrum and the anisotropic spectral eddy viscosity are deduced with an explicit dependence on the anisotropic rotation/stratification time scale which depends on the vertical aspect ratio parameter. Intermediate asymptotic regime corresponding to strong stratification and weak rotation is analyzed where the effects of weak rotation are accounted for by an asymptotic expansion with full control (saturation) of vertical shearing. The regularizing effect of weak rotation differs from regularizations based on vertical viscosity. Two scalar prognostic equations for ageostrophic components (divergent velocity potential and geostrophic departure) are obtained.

del Rosario, R.C.H., and R.C. Smith: *LQR control of shell vibrations via piezoceramic actuators*. ICASE Report No. 97-19, (NASA CR-201673), March 31, 1997, 20 pages. To appear in the Proceedings of the 7th International Conference on Control and Estimation of Distributed Parameter Systems.

A model-based LQR method for controlling vibrations in cylindrical shells is presented. Surface-mounted piezoceramic patches are employed as actuators which leads to unbounded control input operators. Modified Donnell-Mushtari shell equations incorporating strong or Kelvin-Voigt damping are used to model the system. The model is then abstractly formulated in terms of sesquilinear forms. This provides a framework amenable for proving model well-posedness and convergence of LQR gains using analytic semigroup results combined with LQR theory for unbounded input operators. Finally, numerical examples demonstrating the effectiveness of the method are presented.

Mavriplis, Dimitri: *Adaptive meshing techniques for viscous flow calculations on mixed element unstructured meshes*. ICASE Report No. 97-20, (NASA CR-201675), May 29, 1997, 21 pages. Submitted to Journal for Numerical Methods in Fluids.

An adaptive refinement strategy based on hierarchical element subdivision is formulated and implemented for meshes containing arbitrary mixtures of tetrahedra, hexahedra, prisms and pyramids. Special attention is given to keeping memory overheads as low as possible. This procedure is coupled with an algebraic multigrid flow solver which operates on mixed-element meshes. Inviscid flows as well as viscous flows are computed on adaptively refined tetrahedral, hexahedral, and hybrid meshes. The efficiency of the method is demonstrated by generating an adapted hexahedral mesh containing 3 million vertices on a relatively inexpensive workstation.

Ristorcelli, J.R.: *Fluctuating dilatation rate as an acoustic source*. ICASE Report No. 97-21, (NASA CR-201676), May 27, 1997, 11 pages. Submitted to Physics of Fluids.

Ribner's (1962) dilatational acoustic theory is revisited. A rigorous connection between the fluctuating dilatation rate and the acoustic source field is established; this vindicates Ribner's heuristic contention while indicating additional acoustic source terms in his dilatational acoustic theory. It is also shown that Ribner's acoustic source term is quadrupole. Interesting consequences of the dilatational point of view are indicated. It is found that in the region of vortical fluid motion the dilatation scales as the square of the turbulent Mach number, M_t^2 , and has little to do with the acoustic field; its time rate of change, however, is a portion of the sound generation mechanism. Away from the vortical region the fluid dilatation is an acoustic variable and scales as M_t^4 . The mathematical link established is useful to interpreting direct numerical simulation of aeroacoustical flows in which the dilatation is computed.

Ueng, Shyh-Kuang, K. Sikorski, and Kwan-Liu Ma: *Out-of-core streamline visualization on large unstructured meshes*. ICASE Report No. 97-22, (NASA CR-201699), May 29, 1997, 29 pages. Submitted to IEEE TVCG.

It's advantageous for computational scientists to have the capability to perform interactive visualization on their desktop workstations. For data on large unstructured meshes, this capability is not generally available. In particular, particle tracing on unstructured grids can result in a high percentage of non-contiguous memory accesses and therefore may perform very poorly with virtual memory paging schemes. The alternative of visualizing a lower resolution of the data degrades the original high-resolution calculations.

This paper presents an out-of-core approach for interactive streamline construction on large unstructured tetrahedral meshes containing millions of elements. The out-of-core algorithm uses an octree to partition and restructure the raw data into subsets stored into disk files for fast data retrieval. A memory management policy tailored to the streamline calculations is used such that during the streamline construction only a very small amount of data are brought into the main memory on demand. By carefully scheduling computation and data fetching, the overhead of reading data from the disk is significantly reduced and good memory performance results. This out-of-core algorithm makes possible interactive streamline visualization of large unstructured-grid data sets on a single mid-range workstation with relatively low main-memory capacity: 5-20 megabytes. Our test results also show that this approach is much more efficient than relying on virtual memory and operating system's paging algorithms.

Smith, Ralph: *Hysteresis modeling in magnetostrictive materials via Preisach operators*. ICASE Report No. 97-23, (NASA CR-201695), May 7, 1997, 19 pages. Submitted to Journal of Mathematical Systems, Estimation, and Control.

A phenomenological characterization of hysteresis in magnetostrictive materials is presented. Such hysteresis is due to both the driving magnetic fields and stress relations within the material and is significant throughout most of the drive range of magnetostrictive transducers. An accurate characterization of the hysteresis and material nonlinearities is necessary to fully utilize the actuator/sensor capabilities of the magnetostrictive materials. Such a characterization is made here in the context of generalized Preisach operators. This yields a framework amenable to proving the well-posedness of structural models that incorporate the magnetostrictive transducers. It also provides a natural setting in which to develop practical approximation techniques. An example illustrating this framework in the context of a Timoshenko beam model is presented.

Ciardo, Gianfranco, David Nicol, and Kishor S. Trivedi: *Discrete-event simulation of fluid stochastic petri nets*. ICASE Report No. 97-24, (NASA CR-201688), May 7, 1997, 17 pages. Accepted - PNPM '97, St. Malo, France, June 1997.

The purpose of this paper is to describe a method for simulation of recently introduced fluid stochastic Petri nets. Since such nets result in rather complex set of partial differential equations, numerical solution becomes a formidable task. Because of a mixed, discrete and continuous state space, simulative solution also poses some interesting challenges, which are addressed in the paper.

Hayder, M. Ehtesham, Fang Q. Hu, and M. Yousuff Hussaini: *Towards perfectly absorbing boundary conditions for Euler equations*. ICASE Report No. 97-25, (NASA CR-201689), May 9, 1997, 17 pages. Submitted to AIAA Journal.

In this paper, we examine the effectiveness of absorbing layers as non-reflecting computational boundaries for the Euler equations. The absorbing-layer equations are simply obtained by splitting the governing equations in the coordinate directions and introducing absorption coefficients in each split equation. This methodology is similar to that used by Berenger for the numerical solutions of Maxwell's equations. Specifically, we apply this methodology to three physical problems - shock-vortex interactions, a plane free shear flow and an axisymmetric jet - with emphasis on acoustic wave propagation. Our numerical results indicate that the use of absorbing layers effectively minimizes numerical reflection in all three problems considered.

Duck, Peter W., D. Glenn Lasseigne, and M. Y. Hussaini: *The effect of three-dimensional freestream disturbances on the supersonic flow past a wedge*. ICASE Report No. 97-26, (NASA CR-201698), June 9, 1997, 29 pages. Submitted to Physics of Fluids.

The interaction between a shock wave (attached to a wedge) and small amplitude, three-dimensional disturbances of a uniform, supersonic, freestream flow are investigated. The paper extends the two-dimensional study of Duck et al ¹, through the use of vector potentials, which render the problem tractable by the same techniques as in the two-dimensional case, in particular by expansion of the solution by means of a Fourier-Bessel series, in appropriately chosen coordinates.

Results are presented for specific classes of freestream disturbances, and the study shows conclusively that the shock is stable to all classes of disturbances (i.e. time periodic perturbations to the shock do not

grow downstream), provided the flow downstream of the shock is supersonic (loosely corresponding to the weak shock solution). This is shown from our numerical results and also by asymptotic analysis of the Fourier-Bessel series, valid far downstream of the shock.

Interrante, Victoria, Henry Fuchs, and Stephen Pizer: *Conveying the 3D shape of transparent surfaces via texture*. ICASE Report No. 97-27, (NASA CR-201705), June 16, 1997, 32 pages. Submitted to Visualization '96.

Transparency can be a useful device for depicting multiple overlapping surfaces in a single image. The challenge is to render the transparent surfaces in such a way that their three-dimensional shape can be readily understood and their depth distance from underlying structures clearly perceived.

This paper describes our investigations into the use of sparsely-distributed discrete, opaque texture as an "artistic device" for more explicitly indicating the relative depth of a transparent surface and for communicating the essential features of its 3D shape in an intuitively meaningful and minimally occluding way. The driving application for this work is the visualization of layered surfaces in radiation therapy treatment planning data, and the technique is illustrated on transparent isointensity surfaces of radiation dose.

We describe the perceptual motivation and artistic inspiration for defining a stroke texture that is locally oriented in the direction of greatest normal curvature (and in which individual strokes are of a length proportional to the magnitude of the curvature in the direction they indicate), and discuss several alternative methods for applying this texture to isointensity surfaces defined in a volume.

We propose an experimental paradigm for objectively measuring observers' ability to judge the shape and depth of a layered transparent surface, in the course of a task relevant to the needs of radiotherapy treatment planning, and use this paradigm to evaluate the practical effectiveness of our approach through a controlled observer experiment based on images generated from actual clinical data.

Interrante, Victoria: *Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution*. ICASE Report No. 97-28, (NASA CR-201706), June 16, 1997, 16 pages. Submitted to SIGGRAPH '97.

The three-dimensional shape and relative depth of a smoothly curving layered transparent surface may be communicated particularly effectively when the surface is artistically enhanced with sparsely distributed opaque detail.

This paper describes how the set of principal directions and principal curvatures specified by local geometric operators can be understood to define a natural "flow" over the surface of an object, and can be used to guide the placement of the lines of a stroke texture that seeks to represent 3D shape information in a perceptually intuitive way.

The driving application for this work is the visualization of layered isovalue surfaces in volume data, where the particular identity of an individual surface is not generally known a priori and observers will typically wish to view a variety of different level surfaces from the same distribution, superimposed over underlying opaque structures.

By advecting an evenly distributed set of tiny opaque particles, and the empty space between them, via 3D line integral convolution through the vector field defined by the principal directions and principal curvatures of the level surfaces passing through each gridpoint of a 3D volume, it is possible to generate a

single scan-converted solid stroke texture that may intuitively represent the essential shape information of any level surface in the volume.

To generate longer strokes over more highly curved areas, where the directional information is both most stable and most relevant, and to simultaneously downplay the visual impact of directional information in the flatter regions, one may dynamically redefine the length of the filter kernel according to the magnitude of the maximum principal curvature of the level surface at the point around which it is applied.

Strokes are constrained in narrowness by the resolution of the volume within which the texture is represented, but may be variably widened, at the time of rendering to reflect shading information or any other function defined over the volume data, by adaptively indexing into multiple pre-computed texture volumes obtained from advected particles of different sizes.

Bokhari, Shahid H., Thomas W. Crockett, and David M. Nicol: *Binary dissection: Variants & applications*. ICASE Report No. 97-29, (NASA CR-201716), July 22, 1997, 34 pages. To be submitted to IEEE Computational Sciences and Engineering.

Partitioning is an important issue in a variety of applications. Two examples are domain decomposition for parallel computing and color image quantization. In the former we need to partition a computational task over many processors; in the latter we need to partition a high resolution color space into a small number of representative colors. In both cases, partitioning must be done in a manner that yields good results as defined by an application-specific metric.

Binary dissection is a technique that has been widely used to partition non-uniform domains over parallel computers. It proceeds by recursively partitioning the given domain into two parts, such that each part has approximately equal computational load. The basic dissection algorithm does not consider the perimeter, surface area or aspect ratio of the two sub-regions generated at each step and can thus yield decompositions that have poor communication to computation ratios.

We have developed and implemented several variants of the binary dissection approach that attempt to remedy this limitation, are faster than the basic algorithm, can be applied to a variety of problems, and are amenable to parallelization. We first present the Parametric Binary Dissection (PBD) algorithm, which takes into account volume and surface area when partitioning computational domains for use in parallel computing applications. We then consider another variant, the Fast Adaptive Dissection (FAD) algorithm, which provides rapid spatial partitioning for use in color image quantization.

We describe the performance of PBD and FAD on representative problems and present ways of parallelizing the PBD algorithm on 2- or 3-d meshes and on hypercubes.

Chapman, Barbara, Matthew Haines, Piyush Mehrotra, Hans Zima, and John Van Rosendale: *OPUS: A coordination language for multidisciplinary applications*. ICASE Report No. 97-30, (NASA CR-201707), June 19, 1997, 27 pages. To appear in Scientific Programming.

Data parallel languages, such as High Performance Fortran, can be successfully applied to a wide range of numerical applications. However, many advanced scientific and engineering applications are multidisciplinary and heterogeneous in nature, and thus do not fit well into the data parallel paradigm. In this paper we present Opus, a language designed to fill this gap. The central concept of Opus is a mechanism called Shared Abstractions (SDA). An SDA can be used as a computation server, i.e., a locus of computational activity, or as a data repository for sharing data between asynchronous tasks. SDAs can be internally data

parallel, providing support for the integration of data and task parallelism as well as nested task parallelism. They can thus be used to express multidisciplinary applications in a natural and efficient way. In this paper we describe the features of the language through a series of examples and give an overview of the runtime support required to implement these concepts in parallel and distributed environments.

Fenno, C.C., Jr., A. Bayliss, and L. Maestrello: *Interaction of sound from supersonic jets with nearby structures*. ICASE Report No. 97-31, (NASA CR-201708), June 26, 1997, 25 pages. To appear in Proceedings of the 35th AIAA Aerospace Sciences Meeting, Reno, NV.

A model of sound generated in an ideally expanded supersonic (Mach 2) jet is solved numerically. Two configurations are considered; (i) a free jet and (ii) an installed jet with a nearby array of flexible aircraft type panels. In the later case the panels vibrate in response to loading by sound from the jet and the full coupling between the panels and the jet is considered, accounting for panel response and radiation. The long time behavior of the jet is considered. Results for near field and far field disturbance, the far field pressure and the vibration of and radiation from the panels are presented. Panel response crucially depends on the location of the panels. Panels located upstream of the Mach cone are subject to a low level, nearly continuous spectral excitation and consequently exhibit a low level, relatively continuous spectral response. In contrast, panels located within the Mach cone are subject to a significant loading due to the intense Mach wave radiation of sound and exhibit a large, relatively peaked spectral response centered around the peak frequency of sound radiation. The panels radiate in a similar fashion to the sound in the jet, in particular exhibiting a relatively peaked spectral response at approximately the Mach angle from the bounding wall.

Cockburn, Bernardo, and Chi-Wang Shu: *The local discontinuous Galerkin method for time-dependent convection-diffusion systems*. ICASE Report No. 97-32, (NASA CR-201711), July 22, 1997, 35 pages. Submitted to SIAM Journal of Numerical Analysis.

In this paper, we study the Local Discontinuous Galerkin methods for nonlinear, time-dependent convection-diffusion systems. These methods are an extension of the Runge-Kutta Discontinuous Galerkin methods for purely hyperbolic systems to convection-diffusion systems and share with those methods their high parallelizability, their high-order formal accuracy, and their easy handling of complicated geometries, for convection dominated problems. It is proven that for scalar equations, the Local Discontinuous Galerkin methods are L^2 -stable in the nonlinear case. Moreover, in the linear case, it is shown that if polynomials of degree k are used, the methods are k -th order accurate for general triangulations; although this order of convergence is suboptimal, it is sharp for the LDG methods. Preliminary numerical examples displaying the performance of the method are shown.

Kumfert, Gary, and Alex Pothén: *Two improved algorithms for envelope and wavefront reduction*. ICASE Report No. 97-33, (NASA CR-201714), July 15, 1997, 35 pages. To appear in BIT '97.

Two algorithms for reordering sparse, symmetric matrices or undirected graphs to reduce envelope and wavefront are considered. The first is a combinatorial algorithm introduced by Sloan and further developed by Duff, Reid, and Scott; we describe enhancements to the Sloan algorithm that improve its quality and reduce its run time. Our test problems fall into two classes with differing asymptotic behavior of their envelope parameters as a function of the weights in the Sloan algorithm. We describe an efficient $O(n \log n + m)$ time implementation of the Sloan algorithm, where n is the number of rows (vertices), and m is the number of

nonzeros (edges). On a collection of test problems, the improved Sloan algorithm required, on the average, only twice the time required by the simpler Reverse Cuthill-McKee algorithm while improving the mean square wavefront by a factor of three. The second algorithm is a hybrid that combines a spectral algorithm for envelope and wavefront reduction with a refinement step that uses a modified Sloan algorithm. The hybrid algorithm reduces the envelope size and mean square wavefront obtained from the Sloan algorithm at the cost of greater running times. We illustrate how these reductions translate into tangible benefits for frontal Cholesky factorization and incomplete factorization preconditioning.

Loncaric, Josip: *The pseudo-inverse of the derivative operator in polynomial spectral methods*. ICASE Report No. 97-34, (NASA CR-201715), July 22, 1997, 17 pages. Submitted to Journal of Computational Physics.

The matrix $D - kI$ in polynomial approximations of order N is similar to a large Jordan block which is invertible for nonzero k but extremely sensitive to perturbation. Solving the problem $(D - kI)f = g$ involves similarity transforms whose condition number grows as $N!$, which exceeds typical machine precision for $N > 17$. By using orthogonal projections, we reformulate the problem in terms of Q , the pseudo-inverse of D , and therefore its optimal preconditioner. The matrix Q in commonly used Chebyshev or Legendre representations is a simple tridiagonal matrix and its eigenvalues are small and imaginary. The particular solution of $(I - kQ)f = Qg$ can be found for all real k at high resolutions and low computational cost ($O(N)$ times faster than the commonly used Lanczos tau method). Boundary conditions are applied later by adding a multiple of the known homogeneous solution. In Chebyshev representation, machine precision results are achieved at modest resolution requirements. Multidimensional and higher order differential operators can also take advantage of the simple form of Q by factoring or preconditioning.

Interrante, Victoria, and Chester Grosch: *Strategies for effectively visualizing a 3D flow using volume line integral convolution*. ICASE Report No. 97-35, (NASA CR-201717), July 25, 1997, 11 pages. To appear in the Proceedings of Visualization '97, October 1997.

This paper discusses strategies for effectively portraying 3D flow using volume line integral convolution. Issues include defining an appropriate input texture, clarifying the distinct identities and relative depths of the advected texture elements, and selectively highlighting regions of interest in both the input and output volumes. Apart from offering insights into the greater potential of 3D LIC as a method for effectively representing flow in a volume, a principal contribution of this work is the suggestion of a technique for generating and rendering 3D visibility-impeding "halos" that can help to intuitively indicate the presence of depth discontinuities between contiguous elements in a projection and thereby clarify the 3D spatial organization of elements in the flow. The proposed techniques are applied to the visualization of a hot, supersonic, laminar jet exiting into a colder, subsonic coflow.

Zhou, Ye, W. David McComb, and George Vahala: *Renormalization group (RG) in turbulence: Historical and comparative perspective*. ICASE Report No. 97-36, (NASA CR-201718), August 18, 1997, 60 pages. Submitted to Theoretical and Computational Fluid Dynamics.

The terms *renormalization* and *renormalization group* are explained by reference to various physical systems. The extension of renormalization group to turbulence is then discussed; first as a comprehensive review and second concentrating on the technical details of a few selected approaches. We conclude with a discussion of the relevance and application of renormalization group to turbulence modelling.

Ma, Kwan-Liu, and Thomas W. Crockett: *A scalable parallel cell-projection volume rendering algorithm for three-dimensional unstructured data*. ICASE Report No. 97-37, (NASA CR-201719), August 21, 1997, 21 pages. To appear in the Proceedings of the 1997 Parallel Rendering Symposium (Oct. '97).

Visualizing three-dimensional unstructured data from aerodynamics calculations is challenging because the associated meshes are typically large in size and irregular in both shape and resolution. The goal of this research is to develop a fast, efficient parallel volume rendering algorithm for massively parallel distributed-memory supercomputers consisting of a large number of very powerful processors. We use cell-projection instead of ray-casting to provide maximum flexibility in the data distribution and rendering steps. Effective static load balancing is achieved with a round robin distribution of data cells among the processors. A spatial partitioning tree is used to guide the rendering, optimize the image compositing step, and reduce memory consumption. Communication cost is reduced by buffering messages and by overlapping communication with rendering calculations as much as possible. Tests on the IBM SP2 demonstrate that these strategies provide high rendering rates and good scalability. For a dataset containing half a million tetrahedral cells, we achieve two frames per second for a 400x400-pixel image using 128 processors.

Trosset, Michael W., and Virginia Torczon: *Numerical optimization using computer experiments*. ICASE Report No. 97-38, (NASA CR-201724), August 22, 1997, 16 pages. Submitted to Technometrics.

Engineering design optimization often gives rise to problems in which expensive objective functions are minimized by derivative-free methods. We propose a method for solving such problems that synthesizes ideas from the numerical optimization and computer experiment literatures. Our approach relies on kriging known function values to construct a sequence of surrogate models of the objective function that are used to guide a grid search for a minimizer. Results from numerical experiments on a standard test problem are presented.

Sidilkover, David: *Some approaches towards constructing optimally efficient multigrid solvers for the inviscid flow equations*. ICASE Report No. 97-39, (NASA CR-201725), August 21, 1997, 28 pages. To be submitted to Computers and Fluids.

Some important advances took place during the last several years in the development of genuinely multidimensional upwind schemes for the compressible Euler equations. In particular, a robust, high-resolution genuinely multidimensional scheme which can be used for any of the flow regimes computations was constructed. This paper summarizes briefly these developments and outlines the fundamental advantages of this approach.

Swanson, R.C., R. Radespiel, and E. Turkel: *Comparison of several dissipation algorithms for central difference schemes*. ICASE Report No. 97-40, (NASA CR-201726), August 22, 1997, 28 pages. To appear in the Proceedings of the AIAA 13th CFD Conference (of June 29 - July 2, 1997).

Several algorithms for introducing artificial dissipation into a central difference approximation to the Euler and Navier Stokes equations are considered. The focus of the paper is on the convective upwind and split pressure (CUSP) scheme, which is designed to support single interior point discrete shock waves. This scheme is analyzed and compared in detail with scalar and matrix dissipation (MATD) schemes. Resolution capability is determined by solving subsonic, transonic, and hypersonic flow problems. A finite-volume

discretization and a multistage time-stepping scheme with multigrid are used to compute solutions to the flow equations. Numerical results are also compared with either theoretical solutions or experimental data. For transonic airfoil flows the best accuracy on coarse meshes for aerodynamic coefficients is obtained with a simple MATD scheme.

Pilkey, Deborah F., Kevin P. Roc, and Daniel J. Inman: *Computational issues in damping identification for large scale problems*. ICASE Report No. 97-41, (NASA CR-201727), August 21, 1997, 13 pages. Submitted to the Conference on ASME Design Engineering (Sept. 14-17, 1997).

Damage detection and diagnostic techniques using vibration responses that depend on analytical models provide more information about a structure's integrity than those that are not model based. The drawback of these approaches is that some form of workable model is required. Typically, models of practical structures and their corresponding computational effort are very large. One method of detecting damage in a structure is to measure excess energy dissipation, which can be seen in damping matrices. Calculating damping matrices is important because there is a correspondence between a change in the damping matrices and the health of a structure. The objective of this research is to investigate the numerical problems associated with computing damping matrices using inverse methods.

Two damping identification methods are tested for efficiency in large-scale applications. One is an iterative routine, and the other a least squares method. Numerical simulations have been performed on multiple degree-of-freedom models to test the effectiveness of the algorithm and the usefulness of parallel computation for the problems. High Performance Fortran is used to parallelize the algorithm. Tests were performed using the IBM-SP2 at NASA Ames Research Center.

The least squares method tested incurs high communication costs, which reduces the benefit of high performance computing. This method's memory requirement grows at a very rapid rate meaning that larger problems can quickly exceed available computer memory. The iterative method's memory requirement grows at a much slower pace and is able to handle problems with 500+ degrees of freedom on a single processor. This method benefits from parallelization, and significant speedup can be seen for problems of 100+ degrees-of-freedom.

Girimaji, Sharath S., and S. Balachandar: *Analysis and modeling of buoyancy generated turbulence using numerical data*. ICASE Report No. 97-42, (NASA CR-736), September 10, 1997, 35 pages. To be submitted to Int'l Journal of Heat and Mass Transfer.

Rayleigh-Bernard convection offers a unique flow situation in which buoyancy-generated turbulence can be studied in isolation, free of the complicating influence of shear production of turbulence. The objective of this paper is to examine and model important aspects of buoyancy-generated turbulence using direct numerical simulation (DNS) data of Rayleigh-Bernard convection. In particular, we examine the pressure-strain and pressure temperature-gradient correlations, turbulent transport of Reynolds stress, and assumptions pertaining to algebraic modeling of thermal flux and Reynolds stress.

Cockburn, Bernardo, and Chi-Wang Shu: *The Runge-Kutta discontinuous Galerkin method for conservation laws V: Multidimensional systems*. ICASE Report No. 97-43, (NASA CR-201737), September 10, 1997, 37 pages. Submitted to Journal of Computational Physics.

This is the fifth paper in a series in which we construct and study the so-called Runge-Kutta Discontinuous Galerkin method for numerically solving hyperbolic conservation laws. In this paper, we extend

the method to multidimensional nonlinear systems of conservation laws. The algorithms are described and discussed, including algorithm formulation and practical implementation issues such as the numerical fluxes, quadrature rules, degrees of freedom, and the slope limiters, both in the triangular and the rectangular element cases. Numerical experiments for two dimensional Euler equations of compressible gas dynamics are presented that show the effect of the (formal) order of accuracy and the use of triangles or rectangles, on the quality of the approximation.

Ristorcelli, J.R.: *A closure for the compressibility of the source terms in Lighthill's acoustic analogy*. ICASE Report No. 97-44, (NASA CR-201738), September 10, 1997, 30 pages. Submitted to Journal of Fluid Mechanics.

The compressible nature of the source terms in Lighthill's acoustic analogy can be closed. For weakly compressible flows, in the absence of thermoacoustic effects, the compressibility of the source field is known in terms of solenoidal modes of the vortical flow field. In such flows, the square of the fluctuating Mach number is small and this fact, coupled with the singular nature of the acoustic problem, and the fact that the phase speed of the acoustic sources is the advective speed, is used to formally close the compressible portion of the fluctuating Reynolds stresses. The closure resolves, as expressed in Crow's 1970 paper, the inconsistent incompressible approximation to Lighthill's source term. It is shown that the incompressible approximation to Lighthill's source term, accurate to order of the square of the Mach number, predicts an acoustic field accurate to order Mach number.

Kangro, Urve, and Roy Nicolaides: *Divergence boundary conditions for vector Helmholtz equations with divergence constraints*. ICASE Report No. 97-45, (NASA CR-201739), September 11, 1997, 19 pages. Submitted to Mathematical Modeling and Numerical Analysis.

The idea of replacing a divergence constraint by a divergence boundary condition is investigated. The connections between the formulations are considered in detail. It is shown that the most common methods of using divergence boundary conditions do not always work properly. Necessary and sufficient conditions for the equivalence of the formulations are given.

Benkner, Siegfried, Piyush Mehrotra, John Van Rosendale, and Hans Zima: *High-level management of communication schedules in HPF-like languages*. ICASE Report No. 97-46, (NASA CR-201740), September 11, 1997, 27 pages. To be submitted to Supercomputing '97.

The goal of High Performance Fortran (HPF) is to "address the problems of writing data parallel programs where the distribution of data affects performance", providing the user with a high-level language interface for programming scalable parallel architectures and delegating to the compiler the task of producing an explicitly parallel message-passing program. For some applications, this approach may result in dramatic performance losses. An important example is the inspector/executor paradigm, which HPF uses to support irregular data accesses in parallel loops. In many cases, the compiler does not have sufficient information to decide whether an inspector computation is redundant or needs to be repeated. In such cases, the performance of the whole program may be significantly degraded.

In this paper, we describe an approach to solve this problem through the introduction of constructs allowing explicit manipulation of communication schedules at the HPF language level. The goal is to avoid the use of EXTRINSICS for expressing irregular computation via message-passing primitives, while guaranteeing

essentially the same performance. These language features allow the user to control the reuse of schedules and to specify access patterns that may be used to compute a schedule. They are being implemented as part of the HPF+ language and we report some preliminary performance numbers from this implementation.

Guattery, Stephen: *Graph embedding techniques for bounding condition numbers of incomplete factor preconditioners*. ICASE Report No. 97-47, (NASA CR-201741), September 23, 1997, 17 pages. Submitted to Numerical Linear Algebra with Applications.

We extend graph embedding techniques for bounding the spectral condition number of preconditioned systems involving symmetric, irreducibly diagonally dominant M-matrices to systems where the preconditioner is not diagonally dominant. In particular, this allows us to bound the spectral condition number when the preconditioner is based on an incomplete factorization. We provide a review of previous techniques, describe our extension, and give examples both of a bound for a model problem, and of ways in which our techniques give intuitive way of looking at incomplete factor preconditioners.

Ryaben'kii, V.S.: *Difference potentials and their applications*. ICASE Report No. 97-48, (NASA CR-201742), September 24, 1997, 16 pages.

In this lecture, we introduce the concept of difference potentials with the density from the space of discontinuities or jumps, which extends and generalizes the previous constructions of difference potentials; this new concept is sufficiently universal and at the same time simple. The apparatus of difference potentials constitutes the foundation of the difference potentials method (DPM).

Before considering the actual constructions of difference potentials, we discuss some new opportunities that the DPM provides for computations. This brief introductory discussion (as well as the main part of the lecture itself) has a goal of drawing the attention of the scientific computing research community to the DPM and its applications. Although the construction of difference potentials with the density from the space of jumps presents an independent mathematical interest, the subject of this lecture will seem too abstract without discussing the possible applications in the beginning. Moreover, in the end of the lecture we will give a review of the literature emphasizing some applications of the DPM that have already been implemented in computational practice.

Hesthaven, J.S.: *On the analysis and construction of perfectly matched layers for the linearized Euler equations*. ICASE Report No. 97-49, (NASA CR-201744), September 24, 1997, 26 pages. Submitted to Journal of Computational Physics.

We present a detailed analysis of a recently proposed perfectly matched layer (PML) method for the absorption of acoustic waves. The split set of equations is shown to be only weakly well-posed, and ill-posed under small low order perturbations. This analysis provides the explanation for the stability problems associated with the split field formulation and illustrates why applying a filter has a stabilizing effect.

Utilizing recent results obtained within the context of electromagnetics, we develop strongly well-posed absorbing layers for the linearized Euler equations. The schemes are shown to be perfectly absorbing independent of frequency and angle of incidence of the wave in the case of a non-convecting mean flow. In the general case of a convecting mean flow, a number of techniques is combined to obtain a absorbing layers exhibiting PML-like behavior. The efficacy of the proposed absorbing layers is illustrated through computation of benchmark problems in aero-acoustics.

ICASE COLLOQUIA

April 1, 1997 – September 30, 1997

Name/Affiliation/Title	Date
Jameson, Antony, Stanford University "Essential Elements of Computational Algorithms for Aerodynamic Analysis and Design"	April 10
Knight, John C., University of Virginia "Aerospace Software Systems: Challenges and Opportunities"	April 24
Molvig, Kim, MIT, Exa Corporation "High Reynolds Number External Flow Simulations with Digital Physics"	April 28
Papalambros, Panos, University of Michigan, Ann Arbor "Distributed Cooperative Systems Design"	April 29
Hu, Yu Charlie, Harvard University "Data Parallel Hierarchical N-Body Method for Large Scale N-Body Simulations"	May 2
Godunov, Sergei, Sobolev Institute of Mathematics, Novosibirsk, Russia "The Creation of Godunov's Schemes"	May 5
Godunov, Sergei, Sobolev Institute of Mathematics, Novosibirsk, Russia "Spectrum Splitting for Matrix Operators in the Finite Element Method for Elliptic Problems"	May 6
DeCroix, David S., North Carolina State University "Aircraft Wake Vortices and Atmospheric Turbulence A Large Eddy Simulation Approach to Investigate Their Interaction"	May 12
Jiao, Jia, University of Arkansas at Little Rock "Efficient Parallelization of Irregular and Dynamic Problems"	May 13
Spalart, Philippe, Boeing "The Spalart-Allmaras Turbulence Model"	May 14
Couvillion, Warren C., Mississippi State University "Displaying Vector-and-Scalar Fields as Streamlet-Covered Isosurfaces"	May 22

Name/Affiliation/Title	Date
Weigle, Christopher, Mississippi State University "Triangulation of Multivariate Contours for Visualization"	May 29
Whitfield, David, Mississippi State University "Maneuvering Prediction of Self-Propelled Underwater Vehicles"	June 4
Vidimce, Kiril, and David Banks, Mississippi State University "An Interactive 3D Graphics Tool for Virtual Physics"	June 6
Shen, Han-Wei, MRJ, Inc. "UFLIC: A Line Integral Convolution Algorithm for Visualizing Unsteady Flows"	June 19
Mitchell, William, National Institute of Standards and Technology "Overview of a Parallel Hierarchical Adaptive Multigrid Method"	June 27
Leveson, Nancy G., University of Washington "Research in Software System Safety"	June 27
Birge, John, University of Michigan "Stochastic Programming Models in Design"	July 9
Soemarwoto, Bambang I., Institute of Technology, Bandung, India "Aerodynamic Airfoil Optimization Using the Adjoint Method Based on the Euler and RANS Equations"	July 11
Turkel, Eli, Tel-Aviv University, Israel "Low Speed Preconditioning for the Euler and Navier Stokes Equations"	July 14
Forsyth, Peter, University of Waterloo, Canada "Object Oriented Methods for CFD"	July 24
Nowaczyk, Ronald, Clemson University and ICASE "Methods in Group Decision Making: Implications for the Scientific & Engineering Communities"	July 24
Rossow, Cord-Christian, Institut für Entwurfsaerodynamik, Germany "MEGAFLOW - A German CFD Project"	July 25
Nataf, Frederic, Ecole Polytechnique, Paris, France "A Krylov-Schwarz Method with Ventcell Interface Conditions"	July 28

Name/Affiliation/Title	Date
Ghattas, Omar, Carnegie-Mellon University "Large-Scale PDE-Constrained Optimization on Parallel Computers"	July 29
Turkel, Eli, Tel-Aviv University "High Order Schemes for Acoustics and CEM"	July 31
Cook, Jeremy, Parallab, University of Bergen, Norway "A Framework for Workflow Management of Distributed HPC Applications in a Computational Design Context"	August 7
Rossow, Cord-Christian, Institut für Entwurfsaerodynamik, Germany "On a Least Squares Approximation for the Solution of the Euler and Navier-Stokes Equations"	August 8
Mahadevan, Sankaran, Vanderbilt University "Stochastic Optimization in Structural Engineering Design"	August 11
Patten, James M., University of Virginia "Interactive Volume Visualization Using Java and the World Wide Web"	August 19
Keyes, David, Old Dominion University and ICASE "Additive Schwarz Methods with Nonreflecting Boundary Conditions for the Parallel Computation of Helmholtz Problems"	August 20
Ryaben'kii, V.S., Keldysh Institute of Applied Mathematics, Moscow, Russia "The Difference Potentials Method with Density from the 'Space of Jumps' "	August 20
Cross, Mark, and Steve Johnson, University of Greenwich, London, UK "CAPTools - A Software Environment to Transform FORTRAN77 Codes from Scalar to Parallel Form with Message Passing"	August 21
Guerinoni, Fabio, Virginia State University "Parallel PAB3D: Experiences with a Prototype in MPI"	August 21
Cross, Mark, and Steve Johnson, University of Greenwich, London, UK "Software Tools and Technologies for Computational Modelling of Multi-Physics Phenomena in Three Dimensions on MPP Systems"	August 22
Ryaben'kii, V.S., Keldysh Institute of Applied Mathematics, Moscow, Russia "The Problem of Active Shielding"	August 25

Name/Affiliation/Title	Date
Max, Nelson, Lawrence Livermore National Laboratory "Illumination and Shadows in Tree Canopies"	August 26
Chen, Ya-Chin, Imperial College, London, UK "A Novel Approach for Adaptive Signal Processing"	August 27
Bijl, Hester, Delft University of Technology, The Netherlands "A Unified Method for Computing Incompressible and Compressible Flows"	August 28
Hayder, M.E., ICASE, and S.V. Tsynkov, NASA "Computational Methods for Unbounded Domains"	August 29
Ozturan, Can, Bogazici University, Turkey and Abdelkader Baggag, University of Minnesota "Parallelization of an Explicit Aeroacoustics Code Using Parallel Mesh Environment"	August 29
Ristorcelli, J.R., ICASE "A Closure for the Compressibility of the Source Terms in Lighthill's Acoustics Analogy"	September 2
Bokhari, Shahid, University of Engineering, Lahore, Pakistan "Experience with an Unstructured Mesh Code on the HP/Convex SPP-2000"	September 11
Corliss, George, Marquette University "Interval Optimization Techniques, Applications, Sensitivities, and Uncertainties"	September 12
Benilov, A.Y., Stevens Institute of Technology "Sea Surface Aerodynamic and Drag Reduction Problem"	September 19

ICASE SUMMER ACTIVITIES

The summer program for 1997 included the following visitors:

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Abarbanel, Saul <i>Computer Science</i>	Tel-Aviv University	6/30 - 7/25 8/23 - 9/06
Baggag, Abdelkader <i>Computer Science</i>	University of Minnesota	6/16 - 9/12
Banks, David <i>Computer Science</i>	Mississippi State University	5/19 - 6/06
Banks, H. Thomas <i>Applied & Numerical Math</i>	North Carolina State University	6/02 - 6/13
Bayliss, Alvin <i>Fluid Mechanics</i>	Northwestern University	6/23 - 6/27 7/15 - 7/18 7/22 - 7/25 9/02 - 9/11
Bijl, Hester <i>Applied & Numerical Math</i>	Delft University of Technology, The Netherlands	7/21 - 8/29
Blaisdell, Gregory <i>Fluid Mechanics</i>	Purdue University	5/16 - 6/06
Bokhari, Shahid <i>Computer Science</i>	Pakistan University of Engineering	6/13 - 9/12
Booker, Andrew <i>Applied & Numerical Math</i>	The Boeing Company	7/21 - 7/23
Brandt, Achi <i>Applied & Numerical Math</i>	The Weizmann Institute of Science, Israel	6/30 - 7/25
Canuto, Claudio <i>Applied & Numerical Math</i>	Politecnico di Torino, Italy	8/11 - 8/15

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Chapman, Barbara <i>Computer Science</i>	University of Vienna, Austria	9/15 – 10/03
Chen, Ya-Chin <i>Applied & Numerical Math</i>	Imperial College, London, UK	8/04 – 9/12
Chen, Zhikai <i>Computer Science</i>	Old Dominion University	5/12 – 8/08
Chow, Pao-Liu <i>Fluid Mechanics</i>	Wayne State University	6/16 – 6/20 8/18 – 8/22
Ciardo, Gianfranco <i>Computer Science</i>	The College of William and Mary	7/21 – 8/01
Cook, Jeremy <i>Computer Science</i>	Parallab, University of Bergen, Norway	8/04 – 8/08
Couvillion, Warren <i>Computer Science</i>	Mississippi State University	5/19 – 5/23
Cowan, F. Scott <i>Applied & Numerical Math</i>	Georgia Institute of Technology	6/16 – 8/29
Cox, Dennis <i>Applied & Numerical Math</i>	Rice University	7/21 – 7/23
Criminale, William <i>Fluid Mechanics</i>	University of Washington, Seattle	6/16 – 7/11
Darmofal, David <i>Applied & Numerical Math</i>	Texas A&M University	8/04 – 8/15
Diskin, Boris <i>Applied & Numerical Math</i>	The Weizmann Institute of Science, Israel	6/30 – 7/25
Emeric, Pierre <i>Applied & Numerical Math</i>	North Carolina State University	6/02 – 6/13

VISITOR and AREA OF RESEARCH	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Follett, William <i>Applied & Numerical Math</i>	Rocketdyne Division of Boeing	8/04 - 8/08
Forsyth, Peter <i>Computer Science</i>	University of Waterloo, Canada	7/21 - 7/25
Gendron, Guy <i>Applied & Numerical Math</i>	Universite Laval, Canada	8/11 - 8/22
Ghattas, Omar <i>Applied & Numerical Math</i>	Carnegie-Mellon University	7/21 - 8/01
Gottlieb, David <i>Computer Science</i>	Brown University	7/07 - 7/18 7/28 - 8/08 8/25 - 8/28
Grosch, Chester <i>Fluid Mechanics</i>	Old Dominion University	6/02 - 7/31
Guerinoni, Fabio <i>Computer Science</i>	Virginia State University	6/02 - 8/08
Haftka, Raphael <i>Applied & Numerical Math</i>	University of Florida, Gainesville	7/21 - 7/23
Haines, Matthew <i>Computer Science</i>	University of Wyoming	5/12 - 6/06
Hesthaven, Jan <i>Computer Science</i>	Brown University	8/11 - 8/29
Holt, Maurice <i>Applied & Numerical Math</i>	University of California, Berkeley	10/06 - 10/10
Keyes, David <i>Computer Science</i>	Old Dominion University	6/23 - 6/27 7/21 - 8/08 8/18 - 8/29

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Laure, Erwin <i>Computer Science</i>	University of Vienna, Austria	5/19 – 5/30
Max, Nelson <i>Computer Science</i>	Lawrence Livermore Laboratories	8/25 – 8/29
Michielssen, Eric <i>Computer Science</i>	University of Illinois, Urbana-Champaign	6/30 – 7/11
Mitchell, William <i>Computer Science</i>	National Institute of Standards and Technology	6/16 – 6/27
Mittra, Raj <i>Computer Science</i>	Pennsylvania State University	8/04 – 8/06
Mukherjee, Kunal <i>Computer Science</i>	University of Central Florida	8/20 – 8/22
Nataf, Frederic <i>Applied & Numerical Math</i>	Ecole Polytechnique, France	7/14 – 8/08
Nicol, David <i>Computer Science</i>	Dartmouth College	7/21 – 7/28
Nicolaides, R.A. <i>Computer Science</i>	Carnegie-Mellon University	8/04 – 8/08
Owen, Art <i>Applied & Numerical Math</i>	Stanford University	7/21 – 7/23
Ozturan, Can <i>Computer Science</i>	Bogazici University, Turkey	6/25 – 8/29
Painter, James <i>Computer Science</i>	Los Alamos National Laboratory	6/23 – 6/27
Patera, Anthony <i>Applied & Numerical Math</i>	Massachusetts Institute of Technology	7/21 – 7/23

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Patten, James <i>Computer Science</i>	University of Virginia	5/26 – 8/15
Peraire, Jaime <i>Applied & Numerical Math</i>	Massachusetts Institute of Technology	7/21 – 7/23
Pilkey, Deborah <i>Computer Science</i>	Virginia Polytechnic Institute and State University	5/15 – 5/21 6/30 – 8/22
Poje, Andrew <i>Fluid Mechanics</i>	Brown University	7/14 – 7/18
Pothen, Alex <i>Computer Science</i>	Old Dominion University	5/26 6/20
Roe, Philip <i>Applied & Numerical Math</i>	University of Michigan, Ann Arbor	6/02 6/27
Rossow, Cord <i>Applied & Numerical Math</i>	Institut fur Entwurfsaerodynamik, Germany	7/14 8/15
Ryaben'kii, Viktor <i>Applied & Numerical Math</i>	Keldysh Institute for Applied Mathematics, Russia	8/18 9/15
Shen, Han-Wei <i>Computer Science</i>	MRJ Inc.	6/16 6/20
Shih, Ming-Yun <i>Computer Science</i>	The College of William and Mary	5/12 – 8/15
Shu, Chi-Wang <i>Fluid Mechanics</i>	Brown University	5/05 – 5/16 7/07 – 7/11
Simpson, Timothy <i>Applied & Numerical Math</i>	Georgia Institute of Technology	6/16 – 8/29
Soemarwoto, Bambang <i>Applied & Numerical Math</i>	Institute of Technology Bandung, Indonesia	7/07 – 8/01

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Ta'asan, Shlomo <i>Computer Science</i>	Carnegie-Mellon University	6/16 – 7/11
Ting, Lu <i>Fluid Mechanics</i>	New York University	6/16 – 6/20
Torczon, Virginia <i>Computer Science</i>	The College of William and Mary	5/01 – 5/12 6/30 – 7/11 7/28 – 8/08
Trivedi, Kishor <i>Computer Science</i>	Duke University	7/21 – 8/01
Trosset, Michael <i>Applied & Numerical Math</i>	Rice University	7/21 – 7/23
Turkel, Eli <i>Applied & Numerical Math</i>	Tel-Aviv University, Israel	6/16 – 8/15
van Leer, Bram <i>Applied & Numerical Math</i>	The University of Michigan, Ann Arbor	5/06 – 5/09
Vidimce, Kiril <i>Computer Science</i>	Mississippi State University	6/02 – 6/06
Watson, Layne <i>Applied & Numerical Math</i>	Virginia Polytechnic Institute and State University	7/21 – 7/23
Weigle, Christopher <i>Computer Science</i>	Mississippi State University	5/26 – 5/30
Wittenbrink, Craig <i>Computer Science</i>	Hewlett-Packard Laboratories	8/25 – 8/29
Xu, Kun <i>Applied & Numerical Math</i>	Hong Kong University of Science and Technology	6/16 – 6/27
Zima, Hans <i>Computer Science</i>	University of Vienna, Austria	8/14 – 9/17

OTHER ACTIVITIES

On April 10, 1997, ICASE and NASA LaRC co-sponsored the **Second Theodorsen Lectureship Award**. This award was presented to Antony Jameson of Stanford University. Professor Jameson presented a lecture entitled Essential Elements of Computational Algorithms for Aerodynamic Analysis and Design.

On June 23-26, 1997, ICASE sponsored a **Short Course on Gas Automata, Lattice Boltzmann Equation and Gas Kinetics**. The course focused on the principles, advantages, and areas of research in lattice grid methods for fluid flows. There were 26 attendees, and a formal proceedings will be published.

On August 11-13, 1997, ICASE, NASA LaRC and AFOSR co-sponsored the **Symposium on Modeling Complex Turbulent Flows** at the Radisson Hotel in Hampton, VA. The objective of this symposium was to bring together leading researchers in the field of turbulence modeling to (i) evaluate recent progress in modeling, and (ii) anticipate future modeling requirements and preview future research directions. There were 93 attendees, and a formal proceedings will be published.

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970.
Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shelly M. Johnson, Executive Secretary/Visitor Coordinator

Shannon L. Keeter, Technical Publications Secretary (through April 1997)

Rachel A. Lomas, Payroll and Accounting Clerk

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, Senior System Manager

Bryan K. Hess, Assistant System Manager

Gregory P. Wheeler, System Operator

II. SCIENCE COUNCIL

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Dennis Gannon, Professor, Center for Innovative Computer Applications, Indiana University.

Ashwani Kapila, Professor, Department of Mathematics and Science, Rensselaer Polytechnic Institute.

James P. Kendall, Jet Propulsion Laboratory.

Heinz-Otto Kreiss, Professor, Department of Mathematics, University of California at Los Angeles.

Sanjoy Mitter, Professor of Electrical Engineering, Massachusetts Institute of Technology.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Paul Rubbert, Unit Chief, Boeing Commercial Airplane Group.

Ahmed Sameh, Department Head of Computer Science, University of Minnesota.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 2000)

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1997)

Thomas Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Fluid Mechanics. (January 1994 to August 1997)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to August 2000)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to August 1999)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1996 to August 1999)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling]. (January 1995 to July 1997)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to August 1997)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1997)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to August 1999)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to September 1997)

Stephen Guattery - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to August 1997)

M. Ehtesham Hayder - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Fluid Mechanics [Computational Aeroacoustics]. (September 1995 to September 1997)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to August 1998)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics [Multiresolution Schemes]. (August 1996 to October 1997)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to October 1999)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1999)

VI. VISITING SCIENTISTS

Geoffrey M. Lilley - Ph.D., Engineering, Imperial College, London, England, 1945. Professor Emeritus, Department of Aeronautics and Astronautics of Southampton, United Kingdom. Fluid Mechanics [Acroaoustics]. (April 1997 to September 1997)

Ronald H. Nowaczyk - Ph.D., Experimental Psychology, Miami University, 1977. Professor, Department of Psychology, Clemson University. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (January 1997 to September 1997)

John Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Program Director for New Technologies, Division of Advanced Scientific Computing, National Science Foundation. Computer Science [Parallel Computing]. (July 1994 to March 1999)

VII. SHORT TERM VISITING SCIENTISTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Computer Science. (June 1997 to July 1997)

Gregory A. Blaisdell - Ph.D., Mechanical Engineering, Stanford University, 1991. Assistant Professor, School of Aeronautics & Astronautics, Purdue University. Fluid Mechanics. (May 1997 to June 1997)

Shahid Bokhari - Ph.D., Electrical and Computer Engineering, University of Massachusetts-Amherst, 1978. Professor, Department of Electrical Engineering, University of Engineering & Technology, Lahore, Pakistan. Computer Science [Porting NASA-Related codes to the Tera Computer, the "Complete Exchange" Algorithm, and Basic Research on Parallel Computing]. (June 1997 to September 1997)

Andrew J. Booker - Ph.D., Mathematics, University of Washington, 1986. Senior Statistician, Information and Support Services, The Boeing Company. Approximations and Surrogates. (July 1997)

Achi Brandt - Ph.D., Mathematics, Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Israel. Applied & Numerical Mathematics. (June 1997 to July 1997)

Claudio Canuto - Ph.D., Mathematics, University of Turin, Italy, 1975. Professor, Department of Mathematics, Politecnico di Torino, Italy. Applied & Numerical Mathematics. (August 1997)

Jeremy M. Cook - Ph.D., Instrumentation and Analytical Science, The University of Manchester Institute of Science and Technology, 1986. Senior Research Scientist, Parallab, Department of Informatics, University of Bergen. Computer Science. (August 1997)

Dennis D. Cox - Ph.D., Mathematics, University of Washington-Seattle, 1980. Professor, Department of Statistics, Rice University. Approximations and Surrogates. (July 1997)

William O. Criminale - Ph.D., Aeronautics and Mathematics, University of Alabama, 1955. Professor, Department of Applied Mathematics, University of Washington. Fluid Mechanics. (June 1997 to July 1997)

David L. Darmofal - Ph.D., Aerospace Engineering, Massachusetts Institute of Technology, 1993. Assistant Professor, Department of Aerospace Engineering, Texas A&M University. Applied & Numerical Mathematics. (August 1997)

Pierre R. Emeric - Ph.D., Applied Science, The College of William & Mary, 1995. Visiting Assistant Professor, Department of Mathematics, North Carolina State University. Applied & Numerical Mathematics. (June 1997)

William W. Follett - M.S., Mechanical Engineering, Stanford University, 1987. Member of Technical Staff, Advanced Analysis Department, CFD Technology Center. Rocketdyne Division of Boeing. Applied & Numerical Mathematics. (August 1997)

Peter A. Forsyth - Ph.D., Applied Mathematics, University of Western Ontario, 1979. Professor, Department of Computer Science, University of Waterloo, Canada. Approximations and Surrogates. (July 1997)

Guy Gendron - Ph.D., Engineering Science and Mechanics, Virginia Polytechnic Institute and State University, 1991. Assistant Professor, Department of Civil Engineering, Universite Laval, Quebec City, Canada. Applied & Numerical Mathematics. (August 1997)

Omar Ghattas - Ph.D., Mechanics, Duke University, 1984. Associate Professor, Computational Mechanics Laboratory, Department of Civil Engineering, Carnegie Mellon University. Approximations and Surrogates. (July 1997)

Fabio Guerinoni - Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1991. Assistant Professor, Department of Mathematics, Virginia State University. Computer Science. (June 1997 to August 1997)

Raphael Haftka - Ph.D., Aerospace Engineering, University of California-San Diego, 1971. Professor, Department of Aeronautical and Mechanical Engineering, University of Florida. Approximations and Surrogates. (July 1997)

Ignacio M. Llorente - Ph.D., Computer Science, Complutense University of Madrid, Spain, 1995. Associate Professor, Department of Computer Architecture, Complutense University of Madrid, Spain. Computer Science. (September 1997 to December 1997)

Nelson Max - Ph.D., Mathematics, Harvard University, 1967. Computer Scientist, Lawrence Livermore National Laboratory. Computer Science. (August 1997)

Eric Michielssen - Ph.D., Electrical Engineering, University of Illinois at Urbana-Champaign, 1992. Assistant Professor, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign. Computer Science. (June 1997 to July 1997)

William F. Mitchell - Ph.D., Computer Science, University of Illinois-Urbana, 1988. Computer Scientist, Mathematical and Computational Science Division, National Institute of Standards and Technology. Computer Science. (June 1997)

Frederic Nataf - Ph.D., Mathematics, Ecole Polytechnique, France, 1989. Scientist, Center for Applied Mathematics, Ecole Polytechnique, France. Applied & Numerical Mathematics. (July to August 1997)

Art B. Owen Ph.D., Statistics, Stanford University, 1985. Associate Professor, Statistics Department, Stanford University. Approximations and Surrogates. (July 1997)

Can Ozturan - Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1995. Assistant Professor, Computer Engineering Department, Bogazici University, Turkey. Computer Science. (June 1997 to August 1997)

James S. Painter - Ph.D., Computer Science, University of Washington, 1989. Staff Scientist, Advanced Computing Lab, Los Alamos National Laboratory. Computer Science. (June 1997)

Anthony T. Patera - Ph.D., Applied Mathematics, Massachusetts Institute of Technology, 1982. Professor, Department of Mechanical Engineering, Massachusetts Institute of Technology. Approximations and Surrogates. (July 1997)

Jaime Peraire - Ph.D., Numerical Analysis, University of Wales, 1986. Associate Professor, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology. Approximations and Surrogates. (July 1997)

Andrew Poje - Ph.D., Mechanical Engineering, Cornell University, 1993. Research Fellow, Division of Applied Mathematics, Brown University. Fluid Mechanics. (July 1997)

Philip L. Roe - Ph.D., Aeronautics, University of Cambridge, United Kingdom, 1962. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics. (June 1997)

Cord-Christian Rossow - Ph.D., Aerospace Engineering, Technische Universitat Braunschweig, Germany, 1988. Branch Head, Dr.-Ing, Institute for Design Aerodynamics, DLR, Germany. Applied & Numerical Mathematics. (July to August 1997)

Viktor Ryaben'kii - Ph.D., Stability of Difference Equations, Moscow State University, 1953. Leading Research Scientist, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences and Full Professor, Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology. Applied & Numerical Mathematics. (August to September 1997)

Han-Wei Shen - Ph.D., Computer Science, University of Utah, 1997. Research Scientist, NAS Division, NASA Ames Research Center, MRJ Inc. Computer Science. (June 1997)

Bambang I. Soemarwoto - Ph.D., Aeronautics, Delft University of Technology, The Netherlands, 1996. Academic Staff Member, Department of Aerospace Engineering, Institute of Technology, Bandung, Indonesia. Applied & Numerical Mathematics. (July 1997)

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Visiting Associate Professor, Department of Mathematics, University of Arizona. Approximations and Surrogates. (July 1997)

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics. (June 1997 to August 1997)

Layne T. Watson - Ph.D., Mathematics, University of Michigan, 1974. Professor, Department of Computer Science and Mathematics, Virginia Polytechnic Institute and State University. Approximations and Surrogates. (July 1997)

Craig M. Wittenbrink - Ph.D., Electrical Engineering, University of Washington, 1993. Software Engineer, Visual Computing Department, Hewlett-Packard Laboratories. Computer Science. (August 1997)

Kun Xu - Ph.D., Numerical Mathematics, Columbia University, 1993. Assistant Professor, Department of Mathematics, Hong Kong University of Science and Technology. Applied & Numerical Mathematics. (June 1997)

VIII. ASSOCIATE RESEARCH FELLOW

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

IX. CONSULTANTS

Ivo Babuska - Ph.D., Technical Science, Technical University, Prague, Czechoslovakia, 1951; Mathematics, Academy of Science, Prague, 1956; D.Sc., Mathematics, Academy of Science, Prague, 1960. Robert Trull Chair in Engineering, TICAM, The University of Texas at Austin. Applied & Numerical Mathematics [Finite Element Methods Associated With Structural Engineering]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

David Banks - Ph.D., Computer Science, University of North Carolina, 1993. Assistant Professor, Department of Computer Science, Mississippi State University. Computer Science [Graphics and Visualization]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

Oktay Baysal - Ph.D., Mechanical Engineering, Louisiana State University, 1982. Eminent Scholar and Professor, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Mathematics

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Numerical Methods in Optimal Design and Control]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Pao-Liu Chow - Ph.D., Applied Mathematics and Mechanics, Rensselaer Polytechnic Institute, 1967. Professor, Department of Mathematics, Wayne State University. Fluid Mechanics [Acroacoustics and Noise Control]

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant Professor, The College of William & Mary. Computer Science [Reliability Models]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

Dennis B. Gannon - Ph.D., Mathematics, University of California, Davis, 1974. Professor, Department of Computer Science, Indiana University-Bloomington. Computer Science [Investigation of Algorithms and Programming Techniques for Parallel Computers]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Ford Foundation Professor & Chair, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Acoustics]

Matthew D. Haines - Ph.D., Computer Science, Carnegie Mellon University, 1995. Assistant Professor, Department of Computer Science, University of Wyoming. Computer Science [Parallel Programming Environment and Run Time Systems]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Problems Associated with Computational Electromagnetics]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Instability and Transition]

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Fluid Mechanics [Mathematical Combustion]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden, 1960. Professor, Department of Applied Mathematics, California Institute of Technology, Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Scott T. Leutenegger - Ph.D., Computer Science, University of Wisconsin-Madison, 1990. Assistant Professor, Department of Mathematics and Computer Science, University of Denver. Computer Science [System Software Related to Databases for Scientific Data]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Sanjoy K. Mitter - Ph.D., Electrical Engineering, Imperial College of Science & Technology, London, 1965. Professor of Electrical Engineering, Co-Director, Laboratory for Information and Decision Systems, Director, Center for Intelligent Control Systems, Massachusetts Institute of Technology. Fluid Mechanics [Control Theory]

Raj Mittra - Ph.D., Electrical Engineering, University of Toronto, 1957. Professor, Department of Electrical Engineering, Penn State. Computer Science [Computational Electromagnetics]

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, Dartmouth College. Computer Science [Scientific Computing on Scalable Parallel Computers]

R.A. Nicolaides - Ph.D., Computer Science, University of London, 1972. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Ahmed H. Samet - Ph.D., Civil Engineering, University of Illinois, 1968. Head, William Norris Chair, and Professor, Department of Computer Science, University of Minnesota. Computer Science [Numerical Algorithms]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Analysis and Algorithm Development]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics [Nonlinear Acoustic/Structure Interaction Problems]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Design Optimization]

Kishor Trivedi - Ph.D., Computer Science, University of Illinois-Urbana, 1974. Professor, Department of Electrical Engineering, Duke University. Computer Science [Performance and Reliability Modeling Methods, Tools and Applications]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Computational Fluid Dynamics]

Robert G. Voigt - Ph.D., Mathematics, University of Maryland, 1969. Professor, Computational Science Program, The College of William & Mary. Computer Science [High Performance Computing]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-Solvers on Multi Processor Machines]

X. GRADUATE STUDENTS

Abdelkader Baggag - Department of Computer Science, The University of Minnesota. (September 1995 to Present)

Hester Bijl - Faculty of Technical Mathematics and Informatics, Delft University of Technology. (July 1997 to August 1997)

Ya-Chin Chen - Department of Electrical and Electronic Engineering, Imperial College of Science, Technology, and Medicine, London, United Kingdom. (August 1997 to September 1997)

Zhikai Chen - Department of Computer Science, Old Dominion University. (February 1997 to Present)

F. Scott Cowan - Department of Mechanical Engineering, Georgia Institute of Technology. (June 1997 to August 1997)

David C. Cronk - Department of Computer Science, The College of William & Mary. (August 1993 to Present)

Boris Diskin - Department of Applied Mathematics, The Weizmann Institute of Science. (June 1997 to July 1997)

Nilan Karunaratne - Department of Computer Science, Old Dominion University. (August 1995 to Present)

Dinesh Kaushik - Department of Computer Science, Old Dominion University. (May 1997 to Present)

Gary Kumfert - Department of Computer Science, Old Dominion University. (January 1997 to May 1997)

Erwin Laure - Institute for Software Technology and Parallel Systems, University of Vienna. (May 1997)

Georgia Liu - Department of Computer Science, Old Dominion University. (February 1997 to Present)

Joe L. Manthey - Department of Mathematics, Old Dominion University. (September 1993 to June 1997)

Seth D. Milder - Department of Physics and Astronomy, George Mason University. (September 1997 to Present)

Kunal Mukherjee - Department of Computer Science, University of Central Florida. (August 1997)

Deborah F. Pilkey - Department of Engineering Sciences & Mechanics, Virginia Polytechnic Institute and State University. (October 1995 to Present)

Kevin Roe - Department of Computer Science, The College of William & Mary. (May 1995 to Present)

Ming-Yun Shih - Department of Computer Science, The College of William & Mary. (May 1997 to Present)

Timothy W. Simpson - Department of Mechanical Engineering, Georgia Institute of Technology. (June 1997 to August 1997)

Diann P. Smith - Department of Computer Science, The College of William & Mary. (August 1997 to Present)

Praveen Kumar Vangala - Department of Computer Science, Old Dominion University. (November 1996 to August 1997)

Kiril Vidimce - Department of Computer Science, Mississippi State University. (June 1997)

Darren M. Wah - Department of Physics, The College of William & Mary. (February 1997 to Present)

Christopher C. Weigle - Department of Computer Science, Mississippi State University. (May 1997)

XI. STUDENT ASSISTANTS

James Patten - Department of Computer Science, University of Virginia. (May 1997 to August 1997)

Kathryn Paulson - Graduate of Ferguson High School. (June 1997 to August 1997)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE November 1997	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report. April 1, 1997 through September 30, 1997		5. FUNDING NUMBERS C NAS1-97046 C NAS1-19480 WU 505-90-52-01		
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 403, NASA Langley Research Center Hampton, VA 23681-0001		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-97-206258		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified Unlimited Subject Category 59		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period April 1, 1997 through September 30, 1997.				
14. SUBJECT TERMS applied mathematics; multidisciplinary design optimization; fluid mechanics; turbulence; flow control; acoustics; computer science; system software; systems engineering; parallel algorithms			15. NUMBER OF PAGES 90	
			16. PRICE CODE A05	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	